UNIVERSITÄT MANNHEIM

Essays in Collective Decision Making

vorgelegt von

Felix Jarman

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

an der
Fakultät für Rechtswissenschaft und Volkswirtschaftslehre
Abteilung Volkswirtschaftslehre

Juli 2016

Abteilungssprecher: Prof. Dr. Carsten Trenkler

Referent: Prof. Dr. Hans Peter Grüner

Koreferent: Prof. Dr. Thomas Tröger

Verteidigung: 05.09.2016

Declaration of Authorship

I, FELIX JARMAN, declare that this thesis titled, 'Essays in Collective Decision Making' and the work presented in it are my own. Chapter 2 is joint work with Vincent Meisner and was written in close cooperation. No other sources or means except the ones listed have been employed.

Signed: ________________________________

Date: ________________________________

Acknowledgements

First of all, I thank my supervisor Hans Peter Grüner for his guidance and for his support. I experienced him as a great teacher and as somebody who is always chasing the most interesting and relevant questions. It is in large part due to him, that my interests fell where they fell and it was one of his questions that laid the seed for the first paper of this thesis. I also thank my co-supervisor Thomas Tröger. I learned a lot in his seminars and without his invaluable feedback the aforementioned first paper would never have achieved its present form.

I need to thank Vincent Meisner, who is not only my co-author but who is also a close friend. I met Vincent literally on my second day in Mannheim, almost nine years ago. Now, together, we are the proud parents of our joint paper and I am extremely grateful to him.

I would like to thank Benjamin Balzer, Johannes Schneider, and Chris Wolf, with whom I (and also Vincent) not only discussed research, but also football, chess, and life in general, almost every day over lunch. I am grateful to Johannes Bubeck and Christoph Siemroth for being both good office-mates and friends. Christoph always had cookies and was happy to share them.

I thank many, many people in Mannheim and at other places: for reading my research, for pursuing joint research ideas, for listening to me in seminars, for providing helpful comments and advice, for playing squash with me, or just for being fun to be around. I could not possibly name everybody who deserves my thanks, nevertheless, I want to mention (in alphabetical order): Tobias Etzel, Timo Hoffmann, Volker Nocke, Sebastian Siegloch, André Stenzel, and Philipp Zahn.

Finally, I wish to thank my parents Francis and Kyung, my sister Lucy, and my girlfriend Erika. Thank you for always supporting me.

Contents

1. General Introduction

"Collective Decision Making" is a subfield of microeconomics that investigates how public decision-making processes and public organizations work or alternatively how they should work. This dissertation consists of three self-contained papers related to this subfield. In the broadest terms, these papers investigate public decision making processes, i.e., they ask how a decision that involves several individuals is or should be made. The question is not trivial, in particular if these individuals have potentially different preferences and have private information that is relevant to the decision.

If the question is how the the decision making process should be designed, the answer can be found using the tools of mechanism design. Chapter 2, which is joint work with Vincent Meisner, is such a mechanism design exercise. We derive the mechanism that allows for the optimal procurement of projects by a budget-constrained procurement agency under ex-post constraints.

In Chapter 3, I use game theory to analyze a particular decision-making process: voting. In the context of a committee voting on a reform, I ask whether how the possibility of pre-vote cheap talk communication affects outcomes, when the committee consists of two distinct groups.

Finally, in Chapter 4 I look at voting from an empirical perspective. Using data from local elections in Thuringia, I investigate how a small party winning representation on the municipal council affects the council's policies and the party itself.

In the remainder of this introduction, I give a short overview for each chapter.

Chapter 2

In this chapter, we investigate how to design a specific decision-making process. In particular, we ask ourselves how to organize public procurement. Public procurement is an immensely important part of the economy, amounting to, for example, 16 percent of GDP in the EU.[1]

We have a specific public procurement situation in mind. We consider a procurement agency that has a fixed budget and cannot go over budget. It wants to procure projects from different suppliers who propose one project each. The procurement agency knows how much it values each project but the projects' costs are the suppliers' private information. The agency aims to maximize the aggregate value of implemented projects. We want to design a mechanism for this problem that is as robust as possible. Therefore, we impose that all constraints (budget, incentive, and participation) have to hold ex-post.

Our leading example for a concrete procurement problem is that of a development fund that has a fixed budget to build wells in different areas. It knows how valuable a well is in a given community, but it does not know how much money is necessary to build the well.

If costs were not private information, this problem would be the classical knapsack problem: An individual (development fund) has a knapsack that can carry a certain amount of weight (budget). It wants to maximize the value of items (wells) it carries in this knapsack. However, the total weight of carried items (sum of costs of built wells) cannot exceed the capacity of the knapsack.

With private information, this problem becomes a mechanism design problem. However, given our set of constraints, it cannot be solved with the standard pointwise-optimization techniques. Instead, we derive characteristics that the optimal procurement mechanism must have. Those characteristics imply that the optimal mechanism lies within the class of DA-auctions, a class of mechanisms proposed by Milgrom and Segal (2014). As DA auctions are equivalent to clock auctions, this result implies that our optimal procurement mechanism has a corresponding descending-clock auction.

If projects are ex-ante asymmetric, we find a novel and interesting quantity-quality tradeoff: out of two projects the procurement agency must sometimes implement

[1]Source: http://ec.europa.eu/trade/policy/accessing-markets/public-procurement/

the inferior project in order to get both projects with a high probability. This novel tradeoff occurs because here, in contrast to similar mechanisms, the quantity of procured projects is endogenously determined.

Chapter 3

This chapter is about analyzing voting as a decision-making process from a game-theoretical perspective. In particular, it contributes to the literature on the strategic interaction of cheap talk and voting. In it, I consider a committee that consists of two differently sized groups, each with members that have homogeneous preferences. The committee must vote on whether to implement a reform that impacts all members of the committee. Committee members have private information about the probability of success of the reform.

Thus, voting can be seen as a way to aggregate this private information. However, there are other ways to aggregate information. In particular, individuals on the committee could talk to each other prior to voting. This communication is modeled as cheap talk. The novelty in this model is that the committee consists of two different groups whose members can both talk to each other and to the other group.

In this model, I find that among the implementable social choice functions there always exists one Pareto-dominant social choice. This result allows for the comparison of outcomes under two different voting rules: majority rule and unanimity. Surprisingly, there exist parameters such that both groups, including the minority, prefer majority rule and other parameters such that both groups, including the majority, prefer unanimity.

Chapter 4

In this chapter, I empirically investigate representative democracy as a decision-making process. For each party, there is a threshold of minimum votes that the party would need to get in order to win seats[2] on the municipal council. This threshold is a function of how many votes all other parties obtained and follows from the seat-allocation method, which translates vote shares into the number of seats. If a party obtains a number of votes that is very close to this threshold, it can be considered close to random whether this party made it onto the municipal council or not. In most cases in which the party does make it, it gains its seats from other parties who would have been on the municipal council even if the party close to the threshold would not have made it. Thus, in most cases, a small party narrowly making it onto the municipal council implies that there is one additional party on the council.

Using the quasi-randomness in close cases, I can use a regression discontinuity (RD) design to estimate the effect of having seats on the council on the party itself and the effect of having an additional party on the council on the council's policies. I find that a party that narrowly wins seats is 35 percent less likely to drop out in the subsequent election. This result, however, precludes me from estimating a precise incumbency effect, i.e., the effect of being on the council on the vote share in the next election. Estimating bounds on the effect, I find no evidence of a statistically significant incumbency effect.

On the municipal level, I find evidence of an effect on public spending. One additional party leads to an increase in investment spending of almost 50 percent compared to the average level. This increase is larger, the stronger the additional party changes the overall composition of the municipal council.

[2] I use the plural "seats" as in some cases getting an additional vote can make the difference between winning no seats or several seats. In the elections in my dataset, there is a minimum threshold of 5 percent of the vote to win seats. But, on large councils, 5 percent of the vote translates into more than one seat. Thus parties with a vote share of 5 percent can win several seats but parties with 4.99 percent win none.

2. Ex-post Optimal Knapsack Procurement

with Vincent Meisner

Introduction

We study the problem of a procurer who can spend a fixed budget on any of n available projects which differ in the value the designer derives from them. Projects (agents) have private information about their costs and want to get funding beyond the necessary minimum. The designer's goal is to select an affordable set of maximal aggregate quality. In other words, she faces a mechanism design variant of the knapsack problem with strategic behavior due to informational asymmetries.[1] Essentially, we approach this problem as an "up to possibly n-units" procurement problem with n agents with single-unit supply where demand quantity is determined after observing projects' reports under a budget constraint. The budget constraint, the individual rationality constraints, and the incentive compatibility constraints are imposed ex-post, i.e., for any cost realization, implemented projects are always at least fully compensated, the sum of transfers must not exceed the budget, and truth-telling must be a (weakly) dominant strategy. We find that the optimal mechanism can be implemented with a descending-clock auction with a deferred acceptance rule. Because of a tradeoff between quantity and quality, an optimal price clock may have to stop for a period of time leading to instances in which an inferior project is implemented instead of a superior one.

This framework matches a large range of allocation problems, in which a designer needs to allocate a divisible but fixed capacity among agents. Allocation problems,

[1] The knapsack problem is a classical combinatorial problem, dating as far back as 1897. A set of items is assigned values and weights. The knapsack should be filled with the maximal value, but can carry only up to a given weight. For an overview of the literature on knapsack problems, see Kellerer, Pferschy, and Pisinger (2004).

in which a financial budget constraint represents the fixed capacity, include the procurement of bus lines, bridges, and streets, or the allocation of subsidies or research money. Alternatively, the capacity constraint can represent the payload limit on a freighter or on a space shuttle,[2] or a limited amount of time to be devoted to several tasks. Out of many suitable applications, we employ as our leading example a development fund that desires to distribute money to nonprofit projects with nonmonetary benefits.

Our paper not only helps to understand a class of economically relevant problems, the framework also presents a novel methodological challenge. The ex-post nature of both the participation and the budget constraint precludes the use of standard pointwise optimization techniques à la Myerson (1981). Nonetheless, rewriting the problem involves expressing expected transfers in terms of the allocation function as an auxiliary step. As the designer maximizes expected payoff including residual money, we can employ the procurement analogue of Myerson's notion of "virtual values". However, our results qualitatively translate to a setting in which the designer does not value residual money.

By focusing on strategyproof deterministic mechanisms, we can reduce the problem to finding a set of optimal cutoff functions z_i that, for each project i, map the cost vector of other projects $\mathbf{c}_{-i}$ into a cutoff cost level. Project i is conducted if and only if i's cost report falls weakly below cutoff $z_i(\mathbf{c}_{-i})$ and the corresponding compensation payment for that case equals the cutoff $z_i(\mathbf{c}_{-i})$. In optimum, these cutoff functions implement an allocation rule that exhibits certain properties. First, the optimal allocation rule has substitutes: Given a project is implemented for some cost vector, it is also implemented when, all else being equal, the cost of a rival project is increased. Second, the optimal allocation rule has non-bossy winners: A single project that is implemented cannot affect the allocation without changing its own allocation status. Third, the optimal allocation rule excludes all projects with negative "virtual surplus" from the allocation.

By virtue of these properties, any optimal mechanism has an equivalent deferred acceptance (DA) auction representation as described in Milgrom and Segal (2014). A DA auction is an iterative algorithm that computes the allocation and transfers of an auction mechanism and possesses attractive features with respect to bidders'

[2]Clearly, the capacity of a space shuttle is limited. The problem of optimally allocating the capacity and incentivizing projects to reduce payload is economically relevant, see Ledyard, Porter, and Wessen (2000).

incentives that go beyond dominant-strategy implementability. First, in any DA auction, revealing the type truthfully is an "obviously dominant strategy" as defined by Li (2015).[3] Second, any DA auction is weakly group-strategyproof. In other words, it is impossible for a coalition of projects to coordinate their bidding strategies such that it strictly increases the utility of all projects in the coalition. Third, the dominant strategy equilibrium outcome of any DA auction is the only outcome that survives iterated deletion of dominated strategies in the corresponding full information game with the same allocation rule but where players pay their own bid. Therefore predicting the dominant-strategy equilibrium outcome in a DA auction can be considered robust.

Milgrom and Segal (2014) argue that these properties make DA auctions suitable for many challenging environments such as radio spectrum reallocations. Most importantly, they show that every DA auction can be represented by a descending-clock auction. Among several potential applications, they also consider our budget-constrained procurement setup (Example 5: "Adaptive Scoring for a Budget Constraint"). However, they do not show optimality of the DA auction. To the best of our knowledge, we are the first to do so in a nontrivial setting. Therefore we can strengthen the argument in favor of DA auctions. The techniques established in our paper may be helpful to prove optimality of DA auctions in the other settings mentioned in their paper.

Reducing the set of candidates for optimality to a special kind of DA auction implies that any optimal allocation can be implemented with an appropriately designed descending-clock auction: Any project faces a clock with a continuously decreasing price on it, and indicates whether it is willing to conduct its project at this price. In this auction it is a weakly dominant strategy for any project to exit the auction once the clock price hits the project's cost level. At first, we focus on the case in which all projects are ex-ante symmetric: They have the same value and costs are drawn from the same distribution. Here, we show that it is optimal to rank projects according to their cost and "greenlight" the cheapest ones. In optimum, price clocks run down synchronously and hence projects exit in order of their costs until the budget suffices to pay the current clock price to all remaining active projects.

[3]There does not exist any deviation such that, in any information set in which a deviating action is played, the best-case deviation payoff (against even the most favorable profile of strategies of the other players that is consistent with this information set) is strictly larger than the worst-case payoff from truthful bidding (achieved against the least favorable such strategy profile).

Next, we examine the case of ex-ante asymmetric projects, i.e., costs are drawn from different distributions and/or project values differ. Here, we restrict attention to the two-project case because it conveys the main insights while retaining tractability. In applications, the designer may prefer some projects over others and might have different information over cost distributions. In standard procurement settings, the quantity of units to be procured is not endogenously determined as in our model, but it is exogenously fixed to be some quantity k. It is well known that in k-unit procurement auctions the k projects with the greatest non-negative virtual surpluses are implemented, e.g., Luton and McAfee (1986). In the asymmetric case, the ranking implied by costs and the ranking implied by virtual surpluses do not necessarily coincide. Broadly speaking, the designer discriminates against stochastically stronger projects, and favors projects with higher values. The asymmetry requires that each project faces an individual clock and prices decrease asynchronously. In optimum when quantities are exogenous, the clocks' speed is adjusted such that the virtual surplus of marginal projects is kept equal at all times, see Caillaud and Robert (2005, Proposition 1).

Interestingly, the optimal allocation of this environment does not simply translate into the asymmetric case of our environment. In contrast, projects are not always greenlighted in order of their virtual surpluses. Therefore we cannot adopt the approach of Caillaud and Robert (2005). Instead, the descending-clock implementation of the optimal allocation includes individual clocks stopping at certain times. Here, the quantity-quality tradeoff kicks in: We show that the optimal allocation generically features instances in which out of two rival projects the project with lower virtual surplus is chosen. The reasoning behind this result is that the number of procured units is endogenous. In the asymmetric case, always greenlighting in order of virtual surplus reduces the expected number of greenlighted projects compared to the optimal mechanism. Strategyproofness creates a trade-off between quantity and quality of the procured projects. This discrimination of the stronger project is employed on top of the discrimination due to stochastic domination through the virtual costs.

Clock auctions are generally easy to understand and hard to manipulate. Furthermore, they are less information hungry than, for example, sealed bid auctions. In descending-clock auctions, the designer only learns the private information of those projects that are not greenlighted. In fact, Milgrom and Segal (2014) show that clock auctions are the only strategyproof mechanisms that preserve winners' unconditional privacy: Winners only need to reveal the minimum of their private information that is necessary to prove that they should be winning. These features of clock auctions make them attractive for applications in which there is limited trust between the involved parties. In practice, clock auctions are commonly used to sell fish in Japan and they are often found in the public sector, e.g., when the US Department of the Treasury sells warrant positions.

To the best of our knowledge, this paper is the only one that considers purely ex-post constrained optimal procurement design. Such a restrictive setting can be seen as a "worst-case scenario" for the designer, suiting many economic applications. In our leading example of the development fund, an ex-post budget constraint appears natural as budgets are usually fixed. The nonprofit nature of the projects might prohibit acquiring additional money on the financial market. Information rents are necessary, because a project might want to spend money on extra equipment that is convenient for the project's staff but has no value for the designer. In practice, such incentive problems are often resolved using dominant-strategy implementable mechanisms. In strategyproof mechanisms, agents have no incentive to invest in espionage activities or to hire consultants to avoid misspecification of beliefs. Mainly, dominant strategies are desirable as they are easy to explain and not prone to manipulation. For similar reasons, we restrict attention to deterministic mechanisms. Deterministic mechanisms obviate the need for a credible randomization device and are therefore more easily applicable in practice. Finally, ex-post participation constraints are necessary because projects simply cannot be conducted with insufficient funds, and the designer wants to avoid costly renegotiations when the projects default.

Literature

Even though the knapsack problem has a wide range of economic applications, there are relatively few publications in economics on this issue. Most prominently, Maskin (2002), in his Nancy L. Schwartz memorial lecture, addressed the related problem of the UK government that put aside a fixed fund to encourage firms to reduce their pollution. The government faces n firms that have private marginal cost of abatement θ_i and can commit to reduce x_i units of pollution. To reduce pollution as much as possible, the government pays expected compensation transfers t_i to the firms, who report costs and proposed abatement to maximize $t_i - \theta_i x_i$. For some distributions, Maskin (2002) proposes a mechanism that satisfies an ex-post participation constraint, an ex-post incentive compatibility constraint, and the condition that the budget is not exceeded in expectation. In response to Maskin (2002), Chung and Ely (2002b) look at a more general class of mechanism design problems with budget constraints and translate them into a setting à la Baron and Myerson (1982). Their approach nests Maskin (2002) and also Ensthaler and Giebe (2014a) as special cases. However, Ensthaler and Giebe (2014a) more explicitly derive a constructive solution. In contrast to us, they all consider a soft budget constraint that only requires the sum of expected transfers to be less than the budget. By incorporating the budget constraint into a Lagrangian function and ignoring the monotonicity (incentive) constraint, they find a mechanism that, under the standard regularity condition, indeed is incentive compatible.

In addition, Ensthaler and Giebe (2014a) use AGV-budget-balancing (such as Börgers and Norman, 2009) to obtain a mechanism which is ex-post budget-feasible. However, transformation into a mechanism with an ex-post balanced budget in such a way comes at the cost of sacrificing ex-post individual rationality. Many applications do not allow this constraint to be weakened. For instance, subsidy applicants usually cannot be forced to conduct their proposal when receiving only a small or possibly no subsidy. Alternatively, limited liability justifies insisting on ex-post individual rationality. Because we want both constraints to hold ex-post, we cannot build on their techniques and, thus, we approach the problem by characterizing the optimal allocation rule.

To the best of our knowledge, no paper exists that jointly considers optimal mechanism design under ex-post budget balance and ex-post individual rationality in a procurement setting. Ensthaler and Giebe (2014b) propose a belief-free clock mechanism that coincides with our optimal mechanism in the symmetric case for many parameterizations[4] but differs in the asymmetric case by holding the cost-benefit-ratio equal among projects. By simulating different settings, they conclude that this mechanism outperforms a mechanism used in practice. In contrast to their setting, the mechanism designer in our model values residual money. In Section IV, we discuss the meaning of residual money and find that our main results qualitatively carry to the case where residual money is neglected.

[4]For all parameter constellations such that virtual surplus is always nonnegative.

Because of the appeal of dominant-strategy incentive-compatible (DIC) mechanisms compared to Bayesian incentive-compatible (BIC) mechanisms, many researchers have produced valuable BIC-DIC equivalence results. These results characterize environments in which restricting attention to the more robust incentive criterion comes without loss. Our setup is not contained in these environments. For any BIC mechanism, Mookherjee and Reichelstein (1992) show that one can construct a DIC mechanism implementing the same ex-post allocation rule, whenever this allocation rule is monotone in each coordinate. However, the ex-post transfers of the constructed DIC mechanism are not guaranteed to satisfy ex-post budget balance. More recently, Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013) employ a definition of equivalence in terms of interim expected utilities introduced by Manelli and Vincent (2010). For any BIC mechanism, including the optimal one, they construct a DIC mechanism that yields the same interim expected utilities. Here, the ex-post allocation as well as the ex-post transfers might differ between the two. Therefore a DIC mechanism equivalent to a feasible BIC mechanism might violate the ex-post constraints in our setting.

Our budget-constrained procurement setup with ex-post constraints has received much attention in the computer science literature. Instead of specifying the optimal mechanism, the authors in this literature typically aim to construct allocation algorithms that give good approximation guarantees. In other words, they try to maximize the minimal payoff an algorithm can guarantee compared to the full information knapsack payoff. Apart from the seminal paper by Singer (2010), the works of Dobzinski, Papadimitriou, and Singer (2011) and Chen, Gravin, and Lu (2011) are notable examples of this approach. Anari, Goel, and Nikzad (2014) present a stochastic algorithm and show that it gives the best possible approximation guarantee in the many projects limit in which any individual project's costs are small compared to the budget. While the above papers examine the belief-free case, Bei, Chen, Gravin, and Lu (2012) propose an algorithm for setups in which the designer knows how the private information is distributed.

Other auction theoretic papers featuring "knapsack auctions" deal with a slightly different problem compared to us. Aggarwal and Hartline (2006) consider a setting in which each agent is characterized by his object of commonly known size and a privately known valuation for having his object placed in the auctioneer's knapsack with commonly known capacity. They are looking for the truthful auction that best approximates the optimal full-information monotone pricing rule which maximizes the auctioneer's profit. Mu'Alem and Nisan (2008) cover the case of an auctioneer maximizing social welfare instead. Dütting, Gkatzelis, and Roughgarden (2014) study the performance of DA auctions for knapsack auctions, i.e., they show DA auctions fail to achieve a constant factor approximation of the optimal social welfare in knapsack auctions Dizdar, Gershkov, and Moldovanu (2011) investigate a similar knapsack problem of a profit maximizing auctioneer in a dynamic setting: Agents sequentially arrive over time and are either included in the knapsack immediately or lost forever. Thereby they avoid combinatorial issues, which gives rise to a threshold property of the optimal mechanism. In such knapsack auctions, the mechanism designer maximizes the sum of transfers, and the value only enters the individual projects' payoff while the capacity constraint is imposed on the weight assigned to agents. In our framework, the value is collected by the auctioneer and the capacity constraint is imposed on the sum of transfers. Because of the latter, knapsack auctions and our knapsack procurement auctions are not dual problems

There seems to be no reasonable analogy for our setting to another setting in which the mechanism designer is a similarly constrained seller and the agents are buyers. The literature on group-strategyproof cost-sharing mechanisms, initiated by Moulin (1999), considers the dual of a "surplus-sharing" problem. The crucial difference between this problem and our "budget-sharing" problem is that the agents themselves produce the output to be distributed, while in our case the budget to be distributed is fixed and unrelated to the surplus created by the agents, which is collected by the mechanism designer. Budget-constrained buyers in auctions have been discussed in the literature, e.g., by Che and Gale (1998) or Pai and Vohra (2014). However, these authors study budget-constrained agents whereas in our setting the designer is budget-constrained.

In the following section, we introduce the model. In Section III, we rewrite the problem as a problem of finding the optimal cutoff functions and derive a set of properties that any optimal mechanism must have. Sections III.i and III.ii cover symmetric and asymmetric environments, respectively. We discuss extensions and possible modifications to the model in Section IV. Finally, we conclude in Section V.

Model

We consider a set of n projects $I = \{1, \ldots, n\}$ and one mechanism designer. Each project can be conducted exactly once. The designer gains utility v_i if and only if project $i \in I$ is conducted. We consider projects to be utility maximizing agents. If project i is executed, it incurs cost $c_i \in C_i := [\underline{c}_i, \overline{c}_i]$, where in the following we restrict $\underline{c} = 0$.[5] Let $C := \times_{i \in I} C_i$ and $C_{-i} := \times_{j \in I \setminus \{i\}} C_j$. Let the realization of a cost vector be denoted by $\mathbf{c} \in C$. The costs are the projects' private information and are independently drawn from a distribution F_i. We assume F_i to be continuously differentiable with a strictly positive density f_i on the support. The value of the project v_i and the distribution F_i are common knowledge.

To compensate project i for its cost, the designer pays transfer t_i. A direct mechanism is characterized by $\langle q_i, t_i \rangle$. It is a mapping from the vector of cost reports $\mathbf{c} \in C$ into provision decisions and transfers. We denote the allocation function by $\gamma : C \to \mathcal{P}(I)$, and it maps a cost vector into the set of "greenlighted" projects, an element of the power set of I. Correspondingly, we call $I \setminus \gamma(\mathbf{c})$ the set of "redlighted" projects.

We restrict attention to deterministic mechanisms. This restriction implies that once all cost reports are collected, we know with certainty which project is selected by the mechanism. In other words, the decision of implementation q_i is binary,

$$q_i(\mathbf{c}) = \mathbb{I}(i \in \gamma(\mathbf{c})),$$

where $\mathbb{I}$ denotes an indicator function that is one if the corresponding condition is true and zero otherwise. We employ a revelation-principle argument and without loss of generality only consider direct mechanisms.[6]

[5] The impact of this assumption it discussed in Appendix VI.D.

[6] In general, the revelation principle does not hold when restricting attention to deterministic mechanisms: Deterministic direct mechanisms are unable to replicate mixed strategy equilibria in deterministic indirect mechanisms, as noted by, e.g., Strausz (2003). However, in our setting we do not lose generality. A mixed strategy equilibrium consists of a distribution over pure strategy profiles. Because the mechanism is implementable in dominant strategies any of these pure strategy profiles also constitutes a pure strategy equilibrium, in particular the pure strategy equilibrium associated with the designer's most preferred outcome. Similarly, because the mechanism is ex-post constrained, this outcome is feasible. Therefore, while there are allocations that (in the class of deterministic mechanisms) can only be implemented by indirect mechanisms, the designer's most preferred feasible allocation can truthfully be implemented in a direct mechanism.

Project i's utility u_i is given by its transfer minus the cost it bears,

$$u_i(\mathbf{c}) = t_i(\mathbf{c}) - q_i(\mathbf{c})c_i.$$

The designer derives value v_i from each greenlighted project i while having to pay the sum of transfers. Therefore she wants to maximize the aggregate value of greenlighted projects net of transfers paid. Her (ex-post) utility function u_D implies that, in our setting, the designer values residual money,

$$u_D(\mathbf{c}) = \sum_i \big(q_i(\mathbf{c})v_i - t_i(\mathbf{c})\big). \tag{2.1}$$

We impose an ex-post participation constraint. That is, if i is greenlighted the transfer must be at least as high as the cost,

$$t_i(c_i, \mathbf{c}_{-i}) - q_i(c_i, \mathbf{c}_{-i})c_i \geq 0 \quad \forall i \in I, (c_i, \mathbf{c}_{-i}) \in C. \tag{PC}$$

In addition, the designer has a budget constraint which is "hard" in the sense that she cannot spend more than her budget B for any realization of the cost vector. That is, the designer can never exceed her budget,

$$\sum_i t_i(\mathbf{c}) \leq B \quad \forall \mathbf{c} \in C. \tag{BC}$$

Finally, incentive compatibility has to hold ex-post. Alternatively, we can say that the mechanism has to be implementable in (weakly) dominant strategies[7] or that the mechanism must be strategyproof. Therefore for every realization of the cost vector, project i's truthful report must yield at least as much utility as any possible deviation,

$$\begin{aligned} t_i(c_i, \mathbf{c}_{-i}) - q_i(c_i, \mathbf{c}_{-i})c_i &\geq t_i(\widetilde{c}_i, \mathbf{c}_{-i}) - q_i(\widetilde{c}_i, \mathbf{c}_{-i})c_i \\ &\forall i \in I, \mathbf{c}_{-i} \in C_{-i} \text{ and } c_i, \widetilde{c}_i \in C_i. \end{aligned} \tag{IC}$$

[7]In our private value environment, these two concepts are equivalent in a direct revelation mechanism. In general, however, ex-post incentive compatibility is essentially a generalization of dominant-strategy implementability to interdependent value environments. See Chung and Ely (2002a).

Analysis

We search for the direct mechanism that maximizes the expected utility of the designer and refer to this mechanism as the optimal mechanism. One may think that a natural approach to this problem would be to express the ex-post transfer $t_i(c_i, \mathbf{c}_{-i})$ as a function of the ex-post allocation decision $q_i(c_i, \mathbf{c}_{-i})$, taking $\mathbf{c}_{-i}$ as given, and applying the envelope theorem. In that case, it would be possible to restrict attention to the allocation in order to solve for the optimal mechanism. However, this approach does not reduce the complexity of the problem. The reason is that the ex-post transfers and allocation for one cost vector restrict transfers and allocation for other cost vectors through the budget constraint in a manner much more involved than standard monotonicity. In particular, the budget constraint with the ex-post transfer expressed as a function of the ex-post allocation may be ill-behaved. Therefore we cannot straightforwardly arrive at sufficient conditions using convex optimization.[8]

Instead, we derive a set of properties that every mechanism must inherit to be optimal. In general, we establish these properties by showing that the expected payoff yielded by any feasible mechanism not having one of the properties can be increased by adopting the properties. For some of the following lemmata, we provide the proof for the two-project case in the main text and provide the proof of the general case in the appendix. Our first step is to show that strategyproofness implies that the optimal mechanism has to be a cutoff mechanism.

Lemma 2.1. *The optimal mechanism can be represented by cutoff functions $z_i : C_{-i} \to C_i$, such that project i is greenlighted whenever it reports a cost weakly less than its cutoff,*

$$q_i(c_i, \mathbf{c}_{-i}) = \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i})).$$

The transfer to project i equals its cutoff whenever it is greenlighted and zero otherwise,

$$t_i(c_i, \mathbf{c}_{-i}) = q_i(c_i, \mathbf{c}_{-i}) z_i(\mathbf{c}_{-i}).$$

Proof. For any two cost reports $c_i, c_i' \in C_i$ of project i and for some $\mathbf{c}_{-i} \in C_{-i}$, (IC) implies that if the allocation of i is the same, $q_i(c_i, \mathbf{c}_{-i}) = q_i(c_i', \mathbf{c}_{-i})$, also the transfer has to be the same, $t_i(c_i, \mathbf{c}_{-i}) = t_i(c_i', \mathbf{c}_{-i})$. Otherwise, project i could, as one of the cost types, deviate to the report yielding the higher transfer.

[8] Requiring either the budget or the participation constraint to hold only in expectation would enable us to use the techniques employed by Ensthaler and Giebe (2014a).

Conditional on i's allocation and given any cost reports $\mathbf{c}_{-i}$, the transfer is fixed and does not vary with i's cost report. Hence, given $\mathbf{c}_{-i}$, there can only be two different transfers t_i for project i, one for each allocation status, $t_i^{q_i=1}(\mathbf{c}_{-i})$ and $t_i^{q_i=0}(\mathbf{c}_{-i})$.

Define $z_i(\mathbf{c}_{-i}) := t_i^{q_i=1}(\mathbf{c}_{-i}) - t_i^{q_i=0}(\mathbf{c}_{-i})$. Then, (IC) implies

$$q_i(c_i, \mathbf{c}_{-i}) = \begin{cases} 1 & \text{if } c_i \leq z_i(\mathbf{c}_{-i}) \\ 0 & \text{if } c_i > z_i(\mathbf{c}_{-i}) \end{cases}.$$

Suppose to the contrary that for some realization $\widehat{c}_i < z_i(\mathbf{c}_{-i})$ and some other $\widetilde{c}_i < z_i(\mathbf{c}_{-i})$, $q_i(\widehat{c}_i, \mathbf{c}_{-i}) = 0$ and $q_i(\widetilde{c}_i, \mathbf{c}_{-i}) = 1$. Then, type $\widehat{c}_i$ can profitably deviate to reporting $\widetilde{c}_i$ to ensure the green light which yields a utility increase of $z_i(\mathbf{c}_{-i}) - \widehat{c}_i$. An analogous argument applies for $\widehat{c}_i > z_i(\mathbf{c}_{-i}) > 0$.[9]

The last step is to show that $t_i^{q_i=0}(\mathbf{c}_{-i}) = 0$. This result follows from the mechanism being optimal, i.e., maximizing expected utility of the designer. □

As a direct consequence of dominant-strategy implementability, Lemma 2.1 shows that allocation and transfers are characterized by cutoffs. Project i is greenlighted whenever it reports a cost that lies weakly below the cutoff. Crucially, these cutoffs are functions of the other cost reports $\mathbf{c}_{-i}$. However, the optimal cutoffs remain to be determined. The maximization problem of the designer is given by

$$\begin{gathered} \max_{\{z_i\}_{i\in I}} \mathbb{E}_{\mathbf{c}}\left[\textstyle\sum_i q_i(\mathbf{c})v_i - t_i(\mathbf{c})\right] \\ \text{s.t. (BC)}, \\ q_i(\mathbf{c}) = \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i})) \quad \forall \mathbf{c} \in C, \\ t_i(\mathbf{c}) = \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i}))z_i(\mathbf{c}_{-i}) \quad \forall \mathbf{c} \in C. \end{gathered} \tag{2.2}$$

[9] When $c_i = z_i(\mathbf{c}_{-i})$, (IC) permits both $q_i(c_i, \mathbf{c}_{-i}) = 0$ and $q_i(c_i, \mathbf{c}_{-i}) = 1$. By convention, we assume $q_i(c_i, \mathbf{c}_{-i}) = 1$ in this case. However, writing a mechanism this way precludes the specification of tie-breakers, which might be necessary to conserve budget balance. For example, in a two-project example we would write down the mechanism "greenlight the cheaper project" as $z_1(c_2) = c_2$ and $z_2(c_1) = c_1$. If $c_1 = c_2$ a tie-breaker is needed to select a project. As this is a zero-probability event, the choice of the tie-breaker does not impact the designer's payoff. Similarly, as projects are indifferent, their ex-post utility is unaffected. Therefore we refrain from specifying a tie-breaker and proceed with our analysis as if both projects are greenlighted in these cases.

Here, q_i and t_i are determined by the cutoff function z_i. Incentive compatibility and participation constraints, thus, hold by construction.

The next step towards solving this problem involves applying standard methods introduced by Myerson (1981). Let the conditional expected probability of being greenlighted and the conditional expected transfer be

$$\begin{aligned} Q_i(c_i) &= \mathbb{E}_{\mathbf{c}}[q_i(c_i, \mathbf{c}_{-i})|c_i] \\ \text{and } T_i(c_i) &= \mathbb{E}_{\mathbf{c}}[t_i(c_i, \mathbf{c}_{-i})|c_i]. \end{aligned}$$

The interim incentive compatibility required by Myerson (1981) is weaker than our condition (IC). Consequently, the expected transfer is determined by the allocation, $T_i(c_i) = Q_i(c_i)c_i + \int_{c_i}^{\bar{c}_i} Q_i(x)dx$. The usual monotonicity condition is trivially fulfilled as we are dealing with cutoff mechanisms. This reformulation in turn allows us to rewrite the objective function as a function of the allocation. Substituting into problem (2.2) and integrating by parts yields the following maximization problem,

$$\begin{gathered} \max_{\{z_i\}_{i\in I}} \mathbb{E}_{\mathbf{c}} \left[\sum_i \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i})) \left(v_i - c_i - \tfrac{F_i(c_i)}{f_i(c_i)}\right)\right] \\ \text{s.t.} \\ \sum_{i\in I} \mathbb{I}(c_i \leq z_i(\mathbf{c}_{-i}))z_i(\mathbf{c}_{-i}) \leq B \quad \forall \mathbf{c} \in C. \end{gathered} \tag{2.3}$$

We call $\varphi_i(c_i) := c_i + \frac{F_i(c_i)}{f_i(c_i)}$ the *virtual cost* of project i and $\psi_i(c_i) := v_i - \varphi_i(c_i)$ the *virtual surplus*. Here, φ and ψ are the procurement analogues to standard auction terminology. We can directly see from problem (2.3) that the optimal mechanism maximizes the expected sum of greenlighted virtual surpluses.

Note that constrained optimization by Lagrangian is not straightforward here because of the nondifferentiability of the indicator function. Instead, in the following we derive useful properties of the optimal cutoffs that can be exploited to characterize the optimal mechanism. A cutoff mechanism is by construction monotonic in the following sense:

Definition 2.2. An allocation rule γ is monotonic in costs if $i \in \gamma(c_i, \mathbf{c}_{-i})$ and $c_i' < c_i$ imply $i \in \gamma(c_i', \mathbf{c}_{-i})$ for all $\mathbf{c}_{-i} \in C_{-i}$.

In words, if a project gets greenlighted for some cost vector, it also gets greenlighted when, all else equal, its cost is lower. To proceed, we restrict the class of distributions from which costs can be drawn.

Assumption 1 (Log-concavity). For all i, the cumulative distribution function F_i is log-concave.

This assumption is standard in information economics. It is equivalent to the reverse hazard rate function f/F being a weakly decreasing function or the ratio F/f being weakly increasing. Hence, the standard regularity condition is implied: φ_i is strictly increasing and ψ_i is strictly decreasing. A decreasing reverse hazard rate is the procurement analogue to the assumption of increasing hazard rate functions in seller auction settings.

Regularity ensures that a lower cost c_i translates to a higher virtual surplus $\psi_i(c_i)$. Hence, we can define the following cutoff cost type

$$z_i^{**} := \begin{cases} \psi_i^{-1}(0) & \text{if } \psi_i^{-1}(0) \in C_i \\ \overline{c}_i & \text{otherwise} \end{cases}, \tag{2.4}$$

where regularity implies the invertibility of ψ_i and thus allows for the above definition of z_i^{**}. In the symmetric case, $z_i^{**} = z^{**}$ for all $i \in I$. Let $\boldsymbol{\zeta}^{**}$ be the n-dimensional vector with z_i^{**} as i-th element for all $i \in I$.

Definition 2.3. An allocation rule γ is $\boldsymbol{\zeta}^{**}$-exclusive if, for all $i \in I$, $c_i > z_i^{**}$ implies that $i \notin \gamma(c_i, \mathbf{c}_{-i})$ for all $\mathbf{c}_{-i} \in C_{-i}$.

A cutoff mechanism is $\boldsymbol{\zeta}^{**}$-exclusive if and only if $z_i(\mathbf{c}_{-i}) \leq z_i^{**}$ for all $\mathbf{c}_{-i} \in C_{-i}$ and for all $i \in I$. If the budget sufficed, a designer would want to greenlight all projects with nonnegative virtual surplus. Crucially, the arguments leading to this statement also imply that it is never optimal to greenlight a project with negative virtual surplus.

Lemma 2.4. *The optimal mechanism is $\boldsymbol{\zeta}^{**}$-exclusive. In the trivial case, $\sum z_i^{**} \leq B$, the optimal cutoffs are independent of the cost reports,*

$$z_i(\mathbf{c}_{-i}) = z_i^{**} \quad \forall \mathbf{c}_{-i} \in C_{-i} \text{ and } \forall i \in I.$$

The proof of this lemma is standard and hence omitted. It immediately follows from the rewritten objective function (2.3): Greenlighting a project with negative virtual surplus decreases the designer's payoff and uses part of the budget. Guaranteeing the green light for high-cost types comes at the cost of having to pay higher information rents to all cost types. For the same reason, also a budget-unconstrained designer would implement a $\boldsymbol{\zeta}^{**}$-exclusive mechanism, even when the surplus $v_i - \overline{c}_i$ is positive for all projects. Next, we show that an optimal mechanism possesses the following property:

Definition 2.5. An allocation rule γ has substitutes if $i \in \gamma(\mathbf{c})$ and $c_j' > c_j$ for some $j \neq i$ implies $i \in \gamma(c_j', \mathbf{c}_{-j})$.

That is, if a project gets greenlighted for some cost vector $\mathbf{c}$, it is also greenlighted when, all else equal, another project's cost is increased. This property relates to the cross-monotonicity defined in the cost sharing problem of Moulin and Shenker (2001): an agent's cost share cannot increase when the allocation set expands.

Having in mind a setting with an exogenously determined amount of projects to be procured and without a budget constraint, this property is clearly optimal, because if i is among the projects with the highest virtual surpluses for some cost vector, it is also among them when the cost of some other project j is increased, i.e., when j's virtual surplus is decreased. However, with the budget constraint, this property does not hold in a full-information setting.[10] A cutoff mechanism has substitutes if all functions z_i are weakly increasing in each argument.

Lemma 2.6. *The optimal mechanism has substitutes,*

$$z_i(\widetilde{c}_j, \mathbf{c}_{-i-j}) \geq z_i(\widehat{c}_j, \mathbf{c}_{-i-j}) \text{ for almost every } \widetilde{c}_j > \widehat{c}_j \text{ and } \mathbf{c}_{-i-j} \in C_{-i-j}. \quad (2.5)$$

Proof. (with $n = 2$, see appendix for the general proof)
For a graphical representation of the proof, consult Figure 2.1. We show that for any feasible cutoff mechanism that does not have substitutes, there exists a feasible alternative mechanism with substitutes that outperforms the initial candidate in terms of the designer's payoff. In fact, the alternative mechanism outperforms the initial candidate state-by-state and not only in expected terms.

As a first step, we can, without loss of generality, restrict the range of any optimal function z_i: By $\boldsymbol{\zeta}^{**}$-exclusivity, any optimal functional value $z_i(\mathbf{c}_{-i})$ cannot exceed

[10] For example, there are two projects, $v_1 > v_2$. Under full information, both projects get implemented for a cost vector $(c_1, c_2) = (B - z, z)$. Then, increasing c_1 would kick project 2 out of the allocation. In contrast, in our asymmetric-information setting where c_2 pins down a cutoff $z_1(c_2)$ for project 1, project 1 instead loses the green light status, when its cost increases while c_2 remains constant.

z_i^{**}.

Next, fix an arbitrary feasible pair of cutoff functions $\{z_1, z_2\}$ as a candidate for optimality. Contrary to (2.5), suppose that z_2 is decreasing on a set with positive Lebesgue-measure. Then, there exist sets $\widetilde{C}_1$ and $\widehat{C}_1$ with positive Lebesgue-measure, such that

$$z_2(\widehat{c}_1) > z_2(\widetilde{c}_1) \text{ for all } \widehat{c}_1 \in \widehat{C}_1, \widetilde{c}_1 \in \widetilde{C}_1,$$

and $\widehat{c}_1 < \widetilde{c}_1$ for all elements of the corresponding sets.

FIGURE 2.1: The alternative cutoff mechanism $\{z_1, z_2'\}$ outperforms the initial candidate $\{z_1, z_2\}$ for all cost vectors in the light gray area and otherwise yields the same allocation.

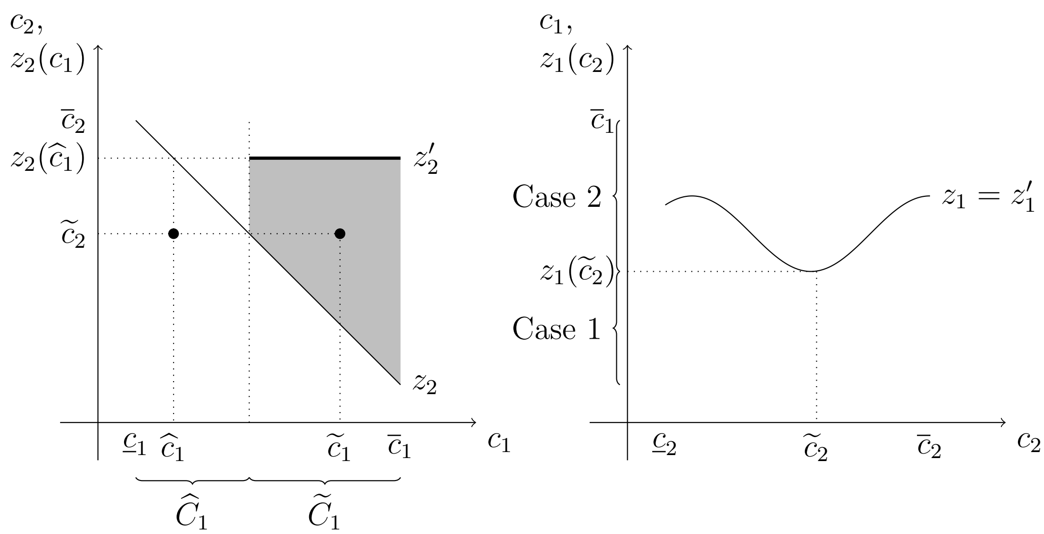

Now, consider an alternative cutoff mechanism $\{z_1, z_2'\}$ that leaves cutoff function z_1 unchanged, but modifies the cutoff function of project 2 in the following way

$$z_2'(c_1) = \begin{cases} z_2(\widehat{c}_1) & \text{if } c_1 \in \widetilde{C}_1 \\ z_2(c_1) & \text{otherwise} \end{cases},$$

with an arbitrary $\widehat{c}_1 \in \widehat{C}_1$. In words, the alternative flattens z_2 over region $\widetilde{C}_1$ and otherwise leaves the initial mechanism as it is. This alternative cutoff function is depicted in Figure 2.1 as the thick flat line.

The alternative mechanism implements the same allocation, except in the gray area depicted in Figure 2.1 where it additionally greenlights project 2. Because

$z_2(\widehat{c}_1) \leq z_2^{**}$ by $\boldsymbol{\zeta}^{**}$-exclusivity, the alternative mechanism clearly yields a higher payoff.

It remains to be shown that the alternative mechanism is not only more profitable but also feasible. First of all, the initial mechanism is, by assumption, budget-feasible everywhere. In particular, it is feasible at any point $(\widehat{c}_1, \widetilde{c}_2)$ with $\widehat{c}_2 \leq z_2(\widehat{c}_1)$ and $\widehat{c}_1 \in \widehat{C}_1$. Formally, for any such points, the budget constraint holds,

$$q_1(\widehat{c}_1, \widetilde{c}_2) z_1(\widetilde{c}_2) + q_2(\widehat{c}_1, \widetilde{c}_2) z_2(\widehat{c}_1) \leq B. \tag{*}$$

To any point $(\widehat{c}_1, \widetilde{c}_2)$, there is a range of corresponding points $(\widetilde{c}_1, \widetilde{c}_2)$ with $\widetilde{c}_1 \in \widetilde{C}_1$. We now check feasibility for any such point $(\widetilde{c}_1, \widetilde{c}_2)$. Referring to Figure 2.1, we are addressing all points that live in the rectangle below the thick flat line of z_2'.

Under the alternative mechanism, for all $\widetilde{c}_1 \in \widetilde{C}_1$, $q_2'(\widetilde{c}_1, \widetilde{c}_2) = q_2(\widehat{c}_1, \widetilde{c}_2) = 1$. Regarding $\widehat{c}_1$, there can be two cases:

Case 1: If $\widehat{c}_1 \leq z_1(\widehat{c}_2)$, then $q_1(\widehat{c}_1, \widetilde{c}_2) = 1$, i.e., both projects are implemented and have to be compensated. The alternative is feasible in any point $(\widetilde{c}_1, \widetilde{c}_2)$ as

$$\begin{aligned} q_1'(\widetilde{c}_1, \widetilde{c}_2) z_1'(\widetilde{c}_2) + q_2'(\widetilde{c}_1, \widetilde{c}_2) z_2'(\widetilde{c}_1) &= q_1'(\widetilde{c}_1, \widetilde{c}_2) z_1(\widetilde{c}_2) + z_2(\widehat{c}_1) \\ &\leq z_1(\widetilde{c}_2) + z_2(\widehat{c}_1) \qquad \leq B, \end{aligned}$$

where the final inequality follows from (*).

Case 2: If $\widehat{c}_1 > z_1(\widehat{c}_2)$, then $q_1(\widehat{c}_1, \widetilde{c}_2) = 0$, i.e., only project 2 is financed. The alternative is feasible in any point $(\widetilde{c}_1, \widetilde{c}_2)$ as

$$\begin{aligned} q_1'(\widetilde{c}_1, \widetilde{c}_2) z_1'(\widetilde{c}_2) + q_2'(\widetilde{c}_1, \widetilde{c}_2) z_2'(\widetilde{c}_1) &= 0 + z_2'(\widetilde{c}_1) \\ &\leq z_2(\widehat{c}_1) \qquad \leq B, \end{aligned}$$

where the first equality follows from $\widetilde{c}_1 \geq \widehat{c}_1 > z_1(\widetilde{c}_2)$ and the final inequality again follows from (*).

Since, for any feasible cutoff mechanism with a cutoff function that is somewhere decreasing, we can find an alternative more profitable cutoff mechanism with cutoff functions that are weakly increasing, the optimal mechanism's allocation rule must have substitutes. $\square$

Lemma 2.6 establishes that optimal cutoff functions are weakly increasing in each of their arguments. The intuition is straightforward. The cost realizations of all projects are independent. Therefore project i's cost report only influences the allocation of project $j \neq i$ via the budget constraint. Project i's cost report only influences the budget through exceeding or lying below the cutoff. If project i exceeds its cutoff, this frees budget to be distributed among the other projects. Consequently, their cutoffs should remain constant or increase. While the intuition is the same for both $n = 2$ and $n > 2$, the proof is more involved in the general case. The reason is that the cost report of the project with the decreasing cutoff does not simultaneously pin down all other cutoffs and the remaining budget - as it does when $n = 2$. We cannot trivially extend the proof above, if some cutoff of a third project z_3 increases in c_1 while z_2 decreases. The intuition of the general proof is that a decreasing cutoff cannot be optimal, because it essentially implies exchanging project 2 for project 1 while the virtual surplus of project 2 decreases relative to the virtual surplus of project 1.

We continue by establishing the next property of the optimal mechanism:

Definition 2.7. An allocation rule γ has non-bossy winners if for any $i \in I$, $\mathbf{c} \in C$, and $c_i' \in C_i$, $i \in \gamma(c_i', \mathbf{c}_{-i}) \cap \gamma(\mathbf{c})$ implies $\gamma(c_i', \mathbf{c}_{-i}) = \gamma(\mathbf{c})$.

In words, a non-bossy winner cannot affect the allocation without changing its own green-light status. In restricted environments, it can be shown that the optimal allocation rule is non-bossy: $\gamma(c_i', \mathbf{c}_{-i}) \cap \{i\} = \gamma(\mathbf{c}) \cap \{i\}$ implies $\gamma(c_i', \mathbf{c}_{-i}) = \gamma(\mathbf{c})$. However, we only need the winners to be non-bossy and examples of environments with bossy losers in the optimal mechanism can be constructed, see Appendix VI.D.

Given some cost vector, let G represent the set of greenlighted projects and R represent the set of redlighted projects. In the following lemma, we show that given that only the projects in some set G are greenlighted and given the remaining projects' costs $\mathbf{c}_R$, for all $g \in G$ all functions z_g intersect at some point $(a_1^G(\mathbf{c}_R), a_2^G(\mathbf{c}_R), ...)$. This point only depends on cost reports $\mathbf{c}_R$ of redlighted projects. Intuitively, optimal cutoffs cannot depend on greenlighted projects' cost, because for these projects the cutoff coincides with the transfer. For the two-project case, Figure 2.2 illustrates that (BC) must bind when both projects are greenlighted. However, then project 1 influencing project 2's cutoff would change the remaining budget which is equal to project 1's transfer, given that (BC) binds. This contradicts the notion of a cutoff mechanism.

Lemma 2.8. *For any cost vectors* $(\mathbf{c}_G, \mathbf{c}_R) \in C$ *and* $(\mathbf{c}'_G, \mathbf{c}_R) \in C$ *such that* $G = \gamma(\mathbf{c}_G, \mathbf{c}_R) = \gamma(\mathbf{c}'_G, \mathbf{c}_R)$ *and* $R = I \setminus \gamma(\mathbf{c}_G, \mathbf{c}_R)$*, the optimal cutoff function* z_g *for all* $g \in G$ *is (almost everywhere) independent of the costs of all greenlighted projects* $\mathbf{c}_G$*. That is,*

$$z_g(\mathbf{c}_{G-g}, \mathbf{c}_R) = z_g(\mathbf{c}'_{G-g}, \mathbf{c}_R),$$

for all $\mathbf{c}_{G-g}$ *and* $\mathbf{c}'_{G-g}$ *such that* G *is the set of greenlighted agents.*

Proof. (with $n = 2$, see appendix for the general proof and consult Figure 2.2 for intuition)
By Lemma 2.1, the optimal mechanism has to be a cutoff mechanism. What remains to be shown is that the cutoff functions $\{z_i\}_{i\in I}$ only depend on $\mathbf{c}_R$. When $\gamma(\mathbf{c})$ is a singleton, i.e., when only one project is greenlighted, the statement follows from the nature of a cutoff function. Hence, we need to show that the cutoffs must be constants whenever $\gamma(\mathbf{c}) = \{1, 2\}$. Therefore suppose that $\gamma(\mathbf{c}) = \{1, 2\}$ is induced with positive probability.

FIGURE 2.2: In Lemma 2.8, we show that in the nontrivial two-project case whenever $G = \{1, 2\}$ both projects get constant transfers summing up to the budget. For instance, the candidate mechanism (with substitutes) depicted above is outperformed by an alternative mechanism indicated by the arrows.

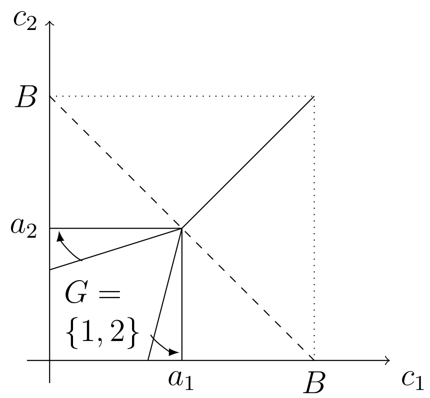

Take any feasible candidate mechanism with any increasing cutoff functions z_i and define

$$\begin{aligned} a_1 &= \max\{c_1 | \exists c_2 : c_2 \le z_2(c_1), c_1 \le z_1(c_2)\} \\ a_2 &= \max\{c_2 | \exists c_1 : c_1 \le z_1(c_2), c_2 \le z_2(c_1)\}, \end{aligned} \tag{2.6}$$

i.e., a_i is the highest cost of project i such that both projects are implemented.

Whenever greenlighting both projects, the sets over which we have defined a_1 and a_2 must be non-empty. The maximum exists by left-continuity of any optimal function z_i.[11] Hence by definition of a_1, there exists $\widetilde{c}_2$ such that $a_1 = z_1(\widetilde{c}_2)$. Similarly, there exists $\widetilde{c}_1$ such that $a_2 = z_2(\widetilde{c}_1)$.

By definition, $(\widetilde{c}_1, \widetilde{c}_2) \leq (a_1, a_2)$ and at cost realization $(\widetilde{c}_1, \widetilde{c}_2)$ both projects are implemented. The budget feasibility of the candidate mechanism implies $a_1 + a_2 \leq B$.

Now we show that, in optimum, $z_1(c_2') = a_1$, for all $c_2' \leq a_2$, and $z_2(c_1') = a_2$, for all $c_1' \leq a_1$. Suppose not. Suppose (without loss of generality) there is some set $\Xi \subset [0, a_2]$ with positive Lebesgue-measure such that $z_1(c_2') < a_1$ for all $c_2' \in \Xi$. Denote $z_1^{\Xi} := \max_{c_2 \in \Xi} z_1(c_2)$. Since $a_1 + a_2 \leq B$, changing the mechanism to $z_1(c_2') = a_1, \quad \forall c_2' \leq a_2$ does not violate the budget constraint and increases the payoff by

$$\Delta > \Pr(c_2 \in \Xi) \int_{z_1^{\Xi}}^{a_1} \psi_1(c) dF(c) > 0.$$

In fact, this alternative mechanism outperforms the initial candidate state-by-state and not only in expectation. $\square$

The following corollary is an immediate consequence of Lemma 2.8 combined with monotonicity and bidder substitutability. It establishes that any optimal mechanism satisfies non-bossiness of greenlighted projects.

Corollary 2.9. *For any optimal mechanism with $G = \gamma(\mathbf{c}_G, \mathbf{c}_R)$ for some $(\mathbf{c}_G, \mathbf{c}_R) \in C$, also $\gamma(\mathbf{c'}_G, \mathbf{c}_R) = G$ for any cost vector $(\mathbf{c'}_G, \mathbf{c}_R) \in C$ with $c_g' \leq c_g$ for all $g \in G$. Hence, for all $i \in I$, for all $\mathbf{c}_{-i} \in C_{-i}$, and for all $\widehat{c}_i$, $\widetilde{c}_i \in C_i$ with $\widehat{c}_i < \widetilde{c}_i$, in any optimal mechanism,*

$$\widehat{c}_i < \widetilde{c}_i \leq z_i(\mathbf{c}_{-i}) \qquad \textit{implies} \qquad \gamma(\widehat{c}_i, \mathbf{c}_{-i}) = \gamma(\widetilde{c}_i, \mathbf{c}_{-i}).$$

Taking stock, among all mechanisms satisfying (PC), (BC) and (IC), any mechanism that maximizes the designer's expected payoff (2.1) belongs to a certain class

[11] We can replace any function z_i with a left-continuous function that is identical up to a set of points with Lebesgue-measure zero. Hence, if there exists an optimal function z_i that is not left-continuous, then there also exists a left-continuous version of the same function that yields the same payoff and hence is also optimal.

of mechanisms: We have shown that the optimal mechanism is characterized by a set of cutoff functions $\{z_i\}_{i\in I}$ and the corresponding allocation rule is

Property 1 monotonic in costs,

Property 2 $\boldsymbol{\zeta}^{**}$-exclusive,

Property 3 has substitutes, and

Property 4 has non-bossy winners.

Being able to restrict attention to mechanisms with these properties is highly useful, as these mechanisms are a much more tangible class than the substantially larger set of all permissible cutoff mechanisms. In addition, all mechanisms with these properties can be implemented with a DA auction as proposed by Milgrom and Segal (2014). To this end, we first restate their definition adapted to our setting.

Definition 2.10 (DA auction)**.** A deferred acceptance (DA) auction is an iterative algorithm defined by a collection of scoring functions

$$s_i^A : C_i \times C_{I\setminus A} \to \mathbb{R}_+$$

that are weakly increasing in c_i for all $i \in A$ and for all $A \subset I$. Let $A_t \subset I$ denote the set of active bidders in iteration t and initially $A_1 = I$. The algorithm stops in some period T when all active projects have a score of zero, $s_i^{A_T} = 0$ for all $i \in A_T$. Then the set of greenlighted project is A_T. Otherwise, at each iteration t, the project with the highest score is removed. The payment p_i^t of project i at iteration t is either given by the highest possible cost that i could have had without being removed from the set of active bidders or by the last iteration's payment, depending on which payment is smaller,

$$p_i^t(\mathbf{c}) = \begin{cases} \sup\{c_i' : s_i^{A_t}(c_i', \mathbf{c}_{I\setminus A_t}) < s_j^{A_t}(c_j, \mathbf{c}_{I\setminus A_t})\} & \text{for } j \in A_t \setminus A_{t+1}, \\ \min\{\sup\{c_i' : s_i^{A_t}(c_i', \mathbf{c}_{I\setminus A_t}) \leq 0\}, p_i^{t-1}\} & \text{if } t = T. \end{cases}$$

The algorithm is initialized with $p_i^0 = \min\{\bar{c}_i, z_i^{**}, B\}$.[12]

The main appeal of DA auctions lies in their incentive guarantees. They are not only strategyproof, they are obviously strategyproof, as defined by Li (2015).

[12] Compared to Milgrom and Segal (2014), we slightly tweak the updating function of payments without changing the deferred acceptance nature of the algorithm and any of its properties.

Moreover, DA auctions are weakly group-strategyproof. That is, no coalition of projects can manipulate their reports such that it strictly increases the utility of all projects in the coalition: At least one member of the coalition receives a weakly worse payoff whenever other coalition members benefit. Because collusion in auctions is generally illegal, compensating the worse off coalition member is not contractible. In addition, the dominant-strategy equilibrium outcome in a DA auction can be interpreted as robust in the following sense: Consider the full-information game in which all cost reports are observed, projects can report any cost, the allocation is determined according to the DA auction's allocation rule, but projects receive their own report as payments. The dominant-strategy equilibrium outcome of the DA auction is the only outcome that survives iterated deletion of dominated strategies in this game.

Proposition 2.11. *Any optimal mechanism has a DA auction representation and can be implemented with a descending-clock auction.*

The proof of Proposition 2.11 is relegated to a separate section in the appendix. Milgrom and Segal (2014) show that with a finite type space, any mechanism satisfying monotonicity, bidder substitutability, and non-bossiness of winners can be implemented by a myopic clock auction.

The symmetric case

In this section, we focus on symmetric projects, i.e., environments with $v_i = v$ and $F_i = F$ for every project $i \in I$. An implication of this assumption is that the order of costs coincides with the order of virtual surpluses and that $z_i^{**} = z^{**}$ for all $i \in I$. We show how to utilize the established results to characterize the optimal allocation and also how to implement it. As in previous proofs, the proof of Proposition 2.12 considers the two-project case while the general proof is relegated to the appendix. In the two-project case, the designer's optimization problem can be reduced to optimally solving for a single constant. Nevertheless, we discuss possible alternatives to the optimal mechanism in greater detail to foreshadow the complications which arise in asymmetric environments.

Proposition 2.12. *Arrange the projects in ascending order of their reported costs,* $c_1 \leq c_2 \leq \cdots \leq c_n \leq c_{n+1} := \overline{c}$, *and define* $z^k := \min\left\{\frac{B}{k}, z^{**}, c_{k+1}\right\}$. *In the symmetric case, the cutoff mechanism with* $z_i(\mathbf{c}_{-i}) = z^{k^*}$ *is the optimal mechanism. The optimal number of accepted projects* k^* *is given by* $k^* := \max\{k | c_k \leq z^k\}$.

Proof. (with $n = 2$, see appendix for the general proof)
In Proposition 2.11, we have shown that the optimal mechanism must be a special kind of DA auction. We call the mechanism in the proposition the proposed mechanism and, as a candidate for optimality, consider $\{z_1, z_2\}$, some different cutoff mechanism with the properties we derived. Suppose $\{z_1, z_2\}$ greenlights both projects with nonzero probability and that it differs from the proposed mechanism in a way such that $a_1 = z > B/2$ and $a_2 = B - z < B/2$ with a_i defined as in (2.6). For graphic intuition of the deviation consult Figure 2.3.

By Lemma 2.4, any optimal mechanism must never greenlight a project with negative virtual surplus. This property is depicted as the kink at (z^{**}, z^{**}).

In the area northwest of the dashed budget line, $c_1 + c_2 > B$, the designer can, by (BC) and (PC), only execute one of the two projects. It can be directly seen from objective function (2.3) that the designer prefers the project with the higher virtual surplus, i.e., the one with lower cost. It does not, however, follow directly that $z_i(c_j) = c_j$ whenever $B - c_i < c_j < z^{**}$. It could be optimal for the designer to forgo executing the lower-cost project for some cost vectors (shaded triangle and crossed square in Figure 2.3) in order to execute both projects in an additional area (horizontally lined, Figure 2.3). In such a case, the designer is forced by incentive compatibility to execute the higher-cost project (for cost vectors in the shaded triangle or the square that is both horizontally and vertically lined).

FIGURE 2.3: A candidate mechanism compared to the proposed mechanism.

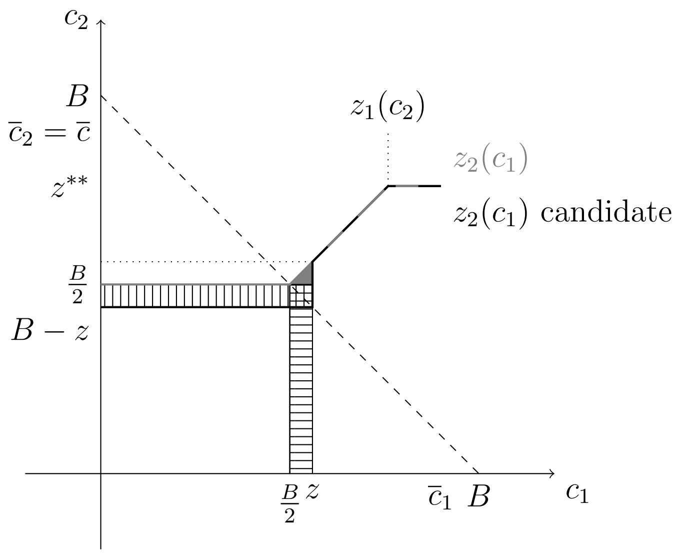

By Lemma 2.8, both cutoffs must be constant whenever both projects are executed. In optimum in that case, there can be no slack in the budget constraint and z_i is flat in that region. Otherwise increasing one of the cutoffs until the budget binds is both feasible and profitable.

Formally, candidate mechanism $\{z_1, z_2\}$ is given by

$$z_2(c_1) = \begin{cases} z^{**} \text{ if } c_1 \geq z^{**} \\ c_1 \text{ if } z < c_1 < z^{**} \\ B - z \text{ if } c_1 < z \end{cases} \quad \text{and} \quad z_1(c_2) = \begin{cases} z^{**} \text{ if } c_2 \geq z^{**} \\ c_2 \text{ if } B - z < c_2 < z^{**} \\ z \text{ if } c_2 < B - z \end{cases} \quad . \tag{2.7}$$

For ease of exposition, let $A = \frac{B}{2}$. Let Δ be the increase in the designer's expected payoff from implementing the proposed mechanism instead of candidate $\{z_1, z_2\}$.

$$\begin{aligned} \Delta = {} & F(z) \int_{B-z}^{A} \psi(x_2) dF(x_2) && \text{(vertical)} \\ & - F(A) \int_{A}^{z} \psi(x_1) dF(x_1) && \text{(horizontal)} \\ & + \int_{A}^{z} \int_{A}^{c} \psi(x_2) dF(x_2) - (F(c) - F(A)) \psi(x_1) dF(x_1) && \text{(shaded)} \end{aligned}$$

where the patterns represent the area in Figure 2.3 where the allocation changes. Everywhere else the allocation and payoff remain the same.

To rewrite Δ, define $\xi(x) = F(x)(v - x)$ with $\xi'(x) = \psi(x) f(x)$:

$$\begin{aligned} \Delta &= F(z)(\xi(A) - \xi(B - z)) - F(A)(\xi(z) - \xi(A)) \\ &+ F(A)(\xi(z) - \xi(A)) + \int_{A}^{z} \xi(x_1) - \xi(A) - F(x_1)\psi_j(x_1) dF(x_1) \\ &= F(z)(\xi(A) - \xi(B - z)) - F(A)(\xi(z) - \xi(A)) \\ &+ F(A)(\xi(z) - \xi(A)) - \xi(A)(F(z) - F(A)) + \int_{A}^{z} F^2(x_1) dx_1 \end{aligned}$$

because $(\psi(c)F(c) - F(c)(v-c))f(c) = F^2(c)$ and then since $\int_A^z F(x_1)^2 dx_1 > F(A)^2 \int_A^z 1 dx_1$,

$$\begin{aligned}\Delta &> F(z)(\xi(A) - \xi(B-z)) - \xi(A)(F(z) - F(A)) + F(A)^2(z-A)\\ &= F(A)^2(v - A + z - A) - F(z)F(B-z)(v - B + z)\\ &= (v - B + z)(F(A)^2 - F(z)F(B-z))\\ &> 0 \quad \Leftrightarrow \quad F(A)^2 > F(z)F(B-z).\end{aligned}$$

This statement is true under Assumption 1, log-concavity. Maximizing $F(z)F(B-z)$ with respect to z, the first order condition is given by

$$\frac{F(z)}{f(z)} = \frac{F(B-z)}{f(B-z)} \tag{2.8}$$

which is only true at $z = B/2$ since F/f is an increasing function. For the same reason, the left-hand side is greater (less) than the right-hand side for $z > B/2 (< B/2)$ making $z = B/2$ the maximum.

We have assumed that in the optimal mechanism both projects get greenlighted for some cost vectors. It remains to show that the optimal mechanism beats the best mechanism in which at most one project gets greenlighted. The best mechanism that selects at most one project always greenlights the project with higher virtual surplus. Clearly, the proposed mechanism outperforms this mechanism as it also always greenlights the project with higher virtual surplus, and it, additionally, sometimes greenlights a second project with positive virtual surplus. $\square$

To sum up, in the symmetric case, the optimal allocation rule takes a simple form: The cheapest projects are greenlighted and the mechanism greenlights as many projects as the budget allows, while each procured project receives the same compensation. Any project that is redlighted prefers this allocation status over having to conduct the project with the associated compensation.

There are two rationales for greenlighted projects to get the same transfer. First, as shown in the proof of Proposition 2.12, this cutoff rule maximizes the probability of getting as many projects as possible. Dominant-strategy incentive compatibility prevents the budget from being shifted away from projects with low cost reports to projects with high costs. Therefore offering equal cutoffs is the best the designer can do. Second, as seen in (2.3), the rewritten maximization problem of the designer, the expected utility of the designer is given by the sum of virtual surpluses of greenlighted projects. Therefore she wants to greenlight those projects with the highest virtual surpluses. That goal is consistent with offering equal cutoffs to greenlighted projects and excluding those with higher cost. In the optimal allocation, greenlighted projects have higher virtual surplus than those which are not greenlighted. The compatibility of the two goals - get as many projects as possible and get those with the highest virtual surpluses - is a special feature of the symmetric case. It generically fails in the asymmetric case, as we demonstrate in the next section.

FIGURE 2.4: An example of optimal allocations for the symmetric case with $n = 2$.

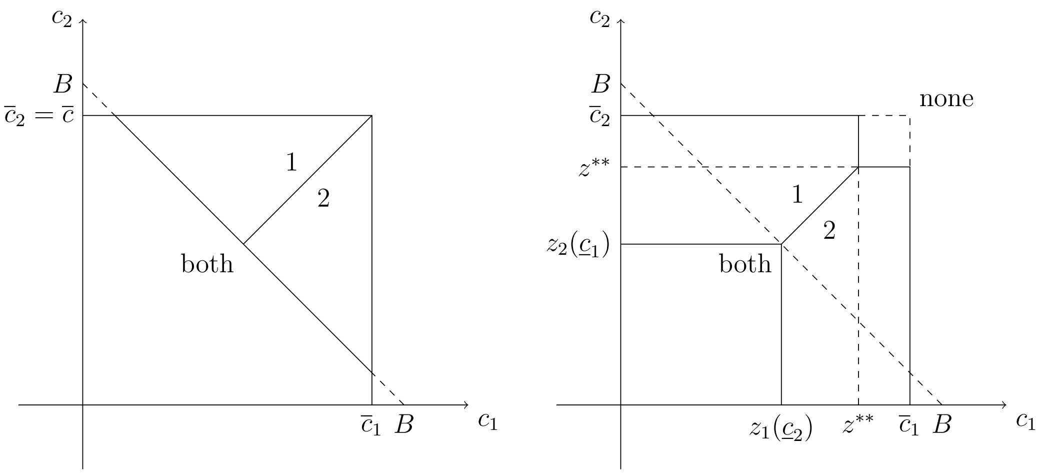

(A) Budget-constrained, full information.

(B) Budget-constrained, private information.

Figure 2.4 illustrates the optimal budget-constrained allocations in an example with two projects. Panel 2.4b shows the fully-constrained optimal allocation juxtaposed with the relaxed optimal allocation when (IC) is neglected, shown in Panel 2.4a. First, note that in this example $v \geq \overline{c}$ and $\overline{c} < B$. Therefore a fully-unconstrained designer with full information would always greenlight both projects, and a budget-constrained designer with full information would always greenlight at least one project. However, since $z^{**} < \overline{c}$, there exist realizations of $\mathbf{c}$ (the upper-right corner of Panel 2.4b) such that no project gets greenlighted in the (IC)-constrained optimal allocation, even though doing so would be profitable from an ex-post perspective. The negative virtual surpluses of the projects in these cases indicates that the cost of allocating to such a project - incentive compatibility requires higher transfers for other cost types - outweighs the benefit from an ex-ante perspective. The second major difference between the relaxed optimal allocation and the optimal allocation can be seen for those realizations of costs such that allocating to both projects would be feasible only in the relaxed problem. This difference is a result of the designer's inability to shift budget from low-cost to relatively higher-cost projects with a strategyproof mechanism.

Corollary 2.13. *In the symmetric case, the optimal direct mechanism can be implemented by a descending-clock auction. The clock price, denoted by τ, starts at z^{**} and descends continuously and synchronously down to $\frac{B}{n}$. Projects can drop out at any price but cannot re-enter. The auction stops once the clock price can be paid out to all projects remaining in the auction.*

In any iteration, a scoring function of the corresponding DA auction is

$$s_i^{A_t}(c_i, A_t) = \max\left\{c_i - \frac{B}{|A_t|}, 0\right\}.$$

We consider the descending-clock auction of Corollary 2.13 to be a natural indirect mechanism that implements the outcome of the optimal allocation. Project i's equilibrium strategy, which implements this outcome, has it staying active as long as the price is weakly larger than its private cost, $\tau \geq c_i$. It is easily verifiable that this is a weakly dominant strategy for project i.

The asymmetric case

In this section, we demonstrate why the logic of the optimal mechanism in the symmetric case does not carry over to the asymmetric case. To preserve tractability, we restrict the analysis to the two-project case which conveys the intuition behind the forces at work in the general case. However, we allow for general values v_1 and v_2 as well as differing cost distributions F_1 and F_2. We consider the non-trivial case, $z_1^{**} + z_2^{**} > B$

Since we did not impose symmetry to prove Proposition 2.11, we can without loss of generality restrict attention to mechanisms inheriting the optimal properties to find an optimal mechanism for the asymmetric case as well. The rewritten maximization problem of the designer (2.3) for the asymmetric two-project case is given by

$$\begin{gathered}\max_{z_1(c_2),z_2(c_1)} \mathbb{E}\Big[\mathbb{I}(c_1 \leq z_1(c_2)) \left(v_1 - c_1 - \tfrac{F_1(c_1)}{f_1(c_1)}\right) \\ +\mathbb{I}(c_2 \leq z_2(c_1)) \left(v_2 - c_2 - \tfrac{F_2(c_2)}{f_2(c_2)}\right) \Big] \\ \text{s.t.} \\ \mathbb{I}(c_1 \leq z_1(c_2))z_1(c_2) + \mathbb{I}(c_2 \leq z_2(c_1))z_2(c_1) \leq B \quad \forall (c_1, c_2) \in C.\end{gathered} \tag{2.9}$$

By Lemma 2.8, the cutoffs must be constants whenever both projects are greenlighted. Since we consider the non-trivial case, these constants must sum up to the budget. Otherwise, increasing one of the cutoffs until the budget binds is both feasible and profitable. Let project 1's cutoff for this case be $z_1(\underline{c}_2) = z$ and project 2's cutoff be $z_2(\underline{c}_1) = B - z$. By virtue of the optimal properties, the designer must greenlight a project once its cost is below the constant cutoff $z_i(\underline{c}_{-i})$. If both projects report greater costs, the designer is free to choose one of them. A glance at the objective function (2.9) reveals that in such a case it is desirable to greenlight the project with greater positive virtual surplus, if feasible. This result allows us to rewrite the objective function (2.9) as a function of z,

$$\begin{aligned}\max_z \pi(z) = &\int_0^z \psi_1(c_1)dF_1(c_1) + \int_0^{B-z} \psi_2(c_2)dF_2(c_2) \\ &+ \int_{\max\{\psi_2^{-1}(\psi_1(z)),B-z\}}^{\bar{c}_2} \int_z^{\min\{\psi_1^{-1}(\psi_2(c_2)),z_1^{**},B\}} \psi_1(x)dF_1(x)dF_2(c_2) \\ &+ \int_{\max\{\psi_1^{-1}(\psi_2(B-z)),z\}}^{\bar{c}_1} \int_{B-z}^{\min\{\psi_2^{-1}(\psi_1(c_1)),z_2^{**},B\}} \psi_2(x)dF_2(x)dF_1(c_1).\end{aligned} \tag{2.10}$$

In the symmetric case, the ranking of virtual surpluses coincides with the reversed order of costs. Hence, the optimal DA auction in the symmetric case rejects in each round the least attractive project in terms of virtual surplus. A natural extension of this mechanism to the asymmetric case would involve adjusting the cutoffs so that they equalize virtual surplus. This modification ensures that again in each round the least attractive project in terms of virtual surplus is rejected. We call this the candidate allocation.

The condition for optimality of the candidate allocation is stated in (2.11). To implement the candidate allocation, the constant cutoffs at which both projects are greenlighted must be a pair $(a_1, a_2) = (z, B - z)$ such that $\psi_1(z) = \psi_2(B - z)$. Then, however, optimality is only obtained if $\frac{F_2(B-z)}{f_2(B-z)} = \frac{F_1(z)}{f_1(z)}$. The intuition behind this statement is straightforward. Selecting z in order to satisfy $\psi_1(z) = \psi_2(B-z)$ allows the designer to always program the price clocks such that they greenlight the project with the higher virtual surplus, whenever it is not feasible to greenlight both projects. However, if $\frac{F_2(B-z)}{f_2(B-z)} \neq \frac{F_1(z)}{f_1(z)}$ the cutoffs z and $B-z$ do not maximize the probability to greenlight both projects. Consequently, the designer can adjust the cutoffs $\{z, B-z\}$ to trade off a higher probability of implementing the most favorable allocation ($\gamma(c_1, c_2) = \{1, 2\}$) against a positive probability of having to implement the less preferred of two possible singleton allocations ($\gamma(\mathbf{c}) = j$, when project j has lower virtual surplus).

Therefore the two aspects of the designer's payoff maximization - getting projects with high virtual surplus and getting as many projects as possible - are only aligned if condition (2.11) is met. In the symmetric case, the condition holds by construction. However, in an asymmetric environment it is generically violated.

Proposition 2.14. *In the nontrivial asymmetric two-project case, i.e., $n = 2$ and $z_1^{**} + z_2^{**} > B$, in which values or cost distributions differ across projects, it is generically not optimal to always greenlight the project with the higher virtual surplus. That is, under the optimal allocation rule γ, there may exist cost vectors $(c_i, c_j, \mathbf{c}_{-i-j}) \in C$ such that*

$$i \notin \gamma(c_i, c_j, \mathbf{c}_{-i-j}), \text{ and } j \in \gamma(c_i, c_j, \mathbf{c}_{-i-j})$$

although

$$\psi_i(c_i) > \psi_j(c_j).$$

Proof. To obtain the derivative of $\pi(z)$ given in (2.10) with respect to z we can use the rules for differentiation under the integral sign.[13] Given the max operators, the derivative takes a different form depending on whether $\psi_1(z) \gtrless \psi_2(B-z)$. However, as π is continuously differentiable, it suffices to look at one of the two forms,

$$\begin{aligned}\left.\frac{\partial \pi}{\partial z}\right|_{z:\psi_1(z)\geq\psi_2(B-z)} &= \int_z^{\psi_1^{-1}(\psi_2(B-z))} \psi_1(x)dF_1(x)f_2(B-z) + \\ &+ \psi_1(z)f_1(z)F_2(B-z) \\ &- \psi_2(B-z)f_2(B-z)F_1(\psi_1^{-1}(\psi_2(B-z))).\end{aligned}$$

[13] Define $g(z, c_2) := \int_z^{\min\{\psi_1^{-1}(\psi_2(c_2)), z_1^{**}, B\}} \psi_1(x)dF_1(x)f_2(c_2)$ and then use $\frac{d}{dz}\left(\int_{a(z)}^{b(z)} g(z, c_2)dc_2\right) = g(z, b(z))b'(z) - g(z, a(z))a'(z) + \int_{a(z)}^{b(z)} g_z(z, c_2)dc_2$.

Now, consider z corresponding to the candidate allocation with $\psi_1(z) = \psi_2(B-z)$, which yields

$$\frac{\partial \pi}{\partial z} = 0 \Leftrightarrow \frac{F_2(B-z)}{f_2(B-z)} = \frac{F_1(z)}{f_1(z)}, \tag{2.11}$$

a nongeneric case. Consequently, it is generically not optimal to always allocate to the project with the higher virtual surplus. □

Proposition 2.14 is driven by a tradeoff between quantity and quality: Even though the designer always prefers the project with the higher virtual surplus, if she greenlights a single project she sometimes greenlights the project with lower virtual surplus out of two rival projects, as quantity is endogenous here. The simplest way to lay out the intuition behind Proposition 2.14 is by an example.

Example 2.1. *There are two projects,* $(n = 2)$ *with* $v_1 = 5, v_2 = 4.5$ *and* c_1 *and* c_2 *are uniformly distributed on support* $[0, 1]$*. The budget is given by* $B = 1$*. The optimal cutoff functions are given by:*

$$z_1(c_2) = \begin{cases} 0.53 & \text{if } c_2 \leq 0.47 \\ c_2 + 0.25 & \text{if } 0.47 < c_2 \leq 0.75 \\ 1 & \text{if } c_2 > 0.75 \end{cases}$$

$$z_2(c_1) = \begin{cases} 0.47 & \text{if } c_1 \leq 0.72 \\ c_1 - 0.25 & \text{if } c_1 > 0.72. \end{cases}$$

Possible scoring functions for a corresponding DA auction are given by:

$$s_1^{\{1,2\}}(c_1) = \begin{cases} c_1 + 0.47 & \text{if } 0.53 < c_1 < 0.72 \\ 2c_1 - 0.25 & \text{if } c_1 \geq 0.72 \\ 0 & \text{otherwise} \end{cases}$$

$$s_2^{\{1,2\}}(c_2) = \begin{cases} 2c_2 + 0.25 & \text{if } c_2 > 0.47 \\ 0 & \text{otherwise} \end{cases}$$

$$s_1^{\{1\}}(c_1) = 0$$

$$s_2^{\{2\}}(c_2) = 0.$$

The corresponding optimal allocation is:

$$(q_1, q_2) = \begin{cases} (1,1) & \text{if } 0 \leq c_1 \leq 0.53 \text{ and } 0 \leq c_2 \leq 0.47 \\ (1,0) & \text{if } 0 \leq c_1 \leq 0.72 \text{ and } c_2 > 0.47 \\ (1,0) & \text{if } c_1 > 0.72 \text{ and } \psi_1 \geq \psi_2 \\ (0,1) & \text{if } 0.53 < c_1 \leq 0.72 \text{ and } c_2 \leq 0.47 \\ (0,1) & \text{if } c_1 > 0.72 \text{ and } \psi_1 < \psi_2. \end{cases}$$

The corresponding transfers are:

$$t_1(c_1, c_2) = \begin{cases} 0.53 & \text{if } c_2 \leq 0.47 \text{ and } c_1 \leq 0.53 \\ c_2 + 0.25 & \text{if } 0.47 < c_2 \leq 0.75 \text{ and } c_1 \leq c_2 + 0.25 \\ 1 & \text{if } c_2 > 0.75 \\ 0 & \text{otherwise} \end{cases}$$

$$t_2(c_1, c_2) = \begin{cases} 0.47 & \text{if } c_1 \leq 0.72 \text{ and } c_2 \leq 0.47 \\ c_1 - 0.25 & \text{if } c_1 > 0.72 \text{ and } c_2 < c_1 - 0.25 \\ 0 & \text{otherwise.} \end{cases}$$

Consider Example 2.1. The candidate allocation demands cutoffs such that $\widetilde{z}_1(\underline{c}_2) = 0.625$ and $\widetilde{z}_2(\underline{c}_1) = 0.375$ for allocating to both projects. At these cutoffs, the probability of greenlighting both projects is $0.625 \cdot 0.375 \approx 0.234$. This allocation is depicted in Panel 2.5a. In contrast, the maximal feasible probability to greenlight both projects is at equal cutoffs, $\widehat{z}_1(\underline{c}_2) = \widehat{z}_2(\underline{c}_1) = 0.5$. The corresponding area is the dotted square in the lower-left corner of Panel 2.5b. However, at these cutoffs it is not incentive compatible to guarantee the green light for the project with higher virtual surplus in every case. More specifically, it is not incentive compatible to allocate along the dotted[14] diagonal line, if at least one project exceeds $\widehat{z}_i(\underline{c}_{-i})$. Hence, strategyproofness introduces a tradeoff between maximizing the probability of greenlighting both projects and allocating to the preferred one if only one project is feasible. Consequently, the optimal cutoffs (z_1^*, z_2^*) for greenlighting

[14]Not to be confused with the dashed diagonal representing the budget constraint.

both projects do not lie at $(0.625, 0.375)$ but rather at $(0.53, 0.47)$. Importantly, this optimal discrimination of the stronger project is pursued independently of the discrimination due to the stochastic dominance reflected in the virtual costs.

Given the optimal allocation in Example 2.1, there are some realizations of the cost vector for which the designer greenlights the project with lower virtual surplus. These realizations are represented by the shaded area in Panel 2.6a. Here, (IC), (PC), and the choice of $(z_1(\underline{c}_2), z_2(\underline{c}_1))$ force the designer to greenlight project 2, even though project 1 has the higher virtual surplus.

The cost vectors for which the designer implements both projects are represented by the rectangular area in the lower-left corner of Panel 2.6a. Any point $(z_1(\underline{c}_2), z_2(\underline{c}_1))$ on the dashed line representing the budget constraint satisfies $z_1(\underline{c}_2)+z_2(\underline{c}_1) = B$. Moving this corner point southwest along the dashed budget line has two effects: shrinking the shaded area and shrinking the area of the rectangle, which in this example represents the probability that both projects are conducted. While it is desirable to shrink the shaded area, in which the designer must allocate to project 2 despite its lower virtual surplus, shrinking the size of the rectangle lowers the probability of allocating to both projects. Given that we have an interior solution in this example, at $(z_1(\underline{c}_2), z_2(\underline{c}_1))$ these two effects balance each other out.

FIGURE 2.5: Candidate and optimal allocation for Example 2.1.

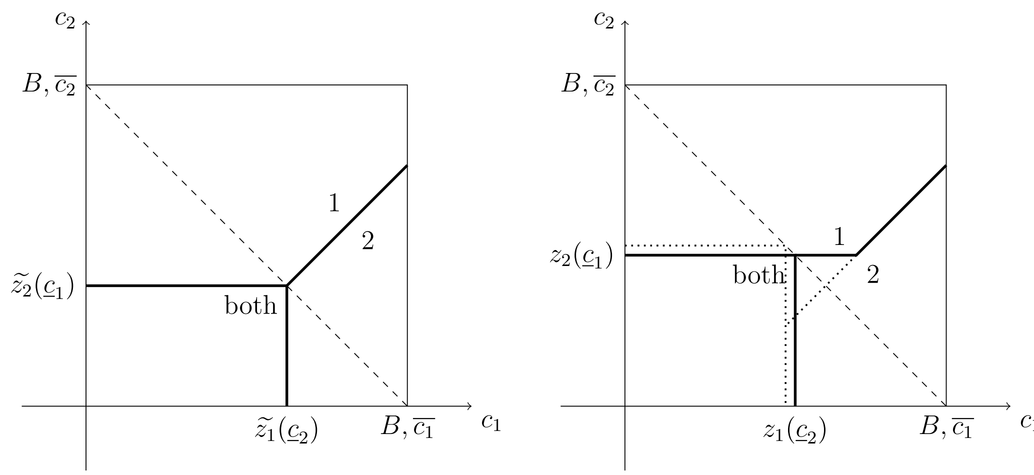

(A) Candidate allocation. (B) Optimal allocation.

Graphically, the fact that there is no slack in the budget constraint whenever both projects are greenlighted implies that the area representing points at which both projects are executed touches the dashed line at least once, as can be seen, for example, in Panel 2.6b. In fact, it can touch the (BC)-constraint exactly once, as it is not possible to greenlight both projects when $c_1 > z_1(\underline{c}_2)$ or $c_2 > z_2(\underline{c}_1)$ without violating (BC) sometimes. This result means that the area where both projects are greenlighted is the rectangle with corners $(0,0)$ and $(z_1(\underline{c}_2), z_2(\underline{c}_1))$. Then, if $c_1 < z_1(\underline{c}_2)$ but $c_2 > z_2(\underline{c}_1)$, the nature of cutoffs prevents the designer from greenlighting project 2. Therefore project 1 must be greenlighted, as represented by the lightly shaded area in Panel 2.6b. A similar argument applies to the darkly shaded area. Thus, looking at Panel 2.6b, the choice of $(z_1(\underline{c}_2), z_2(\underline{c}_1))$ determines the allocation for all cost realizations except those in the upper-right corner. Here, the designer is free to choose the allocation, as long as the line delineating whether project 1 or 2 gets greenlighted is (weakly) increasing or vertical. Not surprisingly, it is optimal to greenlight the project with the higher virtual surplus.

FIGURE 2.6: Greenlighting the project with lower virtual surplus and (IC)-constraints on the allocation (Example 2.1).

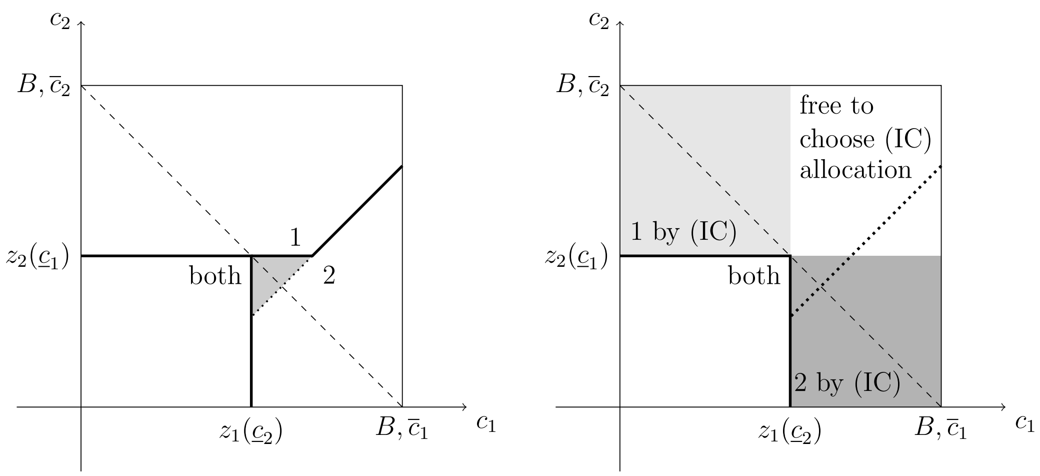

(A) Greenlighting the project with lower virtual surplus.

(B) (IC)-constraints on the allocation.

By Proposition 2.11, the optimal allocation can be implemented with a descending-clock auction. In the following, we show how to accommodate the tradeoff between quantity and quality in a modified clock auction.

Corollary 2.15. *In an optimal implementation with descending price clocks, the clocks not only run at individual speeds, occasionally some clocks also have to halt.*

A crucial difference to the symmetric case is that each project must have an individual price clock, because heterogeneous virtual surplus functions require individual speeds. Interestingly, an implication of the quantity-quality tradeoff is that sometimes one clock has to halt. For Example 2.1, the clock prices, denoted by τ_i, are depicted in Figure 2.7 as a function of time. The entire (maximal) duration of the auction can be divided into three segments. The auction starts with both clocks at $z_1^{**} = z_2^{**} = \bar{c}$. First, τ_2 decreases while τ_1 is held constant, which happens until both clock prices lead to the same virtual surplus, i.e., $\psi_2(\tau_2) = \psi_1(\bar{c}_2)$. Second, both τ_1 and τ_2 decrease simultaneously, but asynchronously keeping virtual surplus equal, $\psi_1(\tau_1) = \psi_2(\tau_2)$, until $\tau_2 = z_2(\underline{c}_1)$. Third, only τ_1 decreases until $\tau_1 = z_1(\underline{c}_2)$. If at this point both projects still remain in the auction, the auction stops and both are greenlighted. Otherwise, the inferior project 2 is greenlighted.

FIGURE 2.7: Optimal descending-clock auction in Example 2.1.

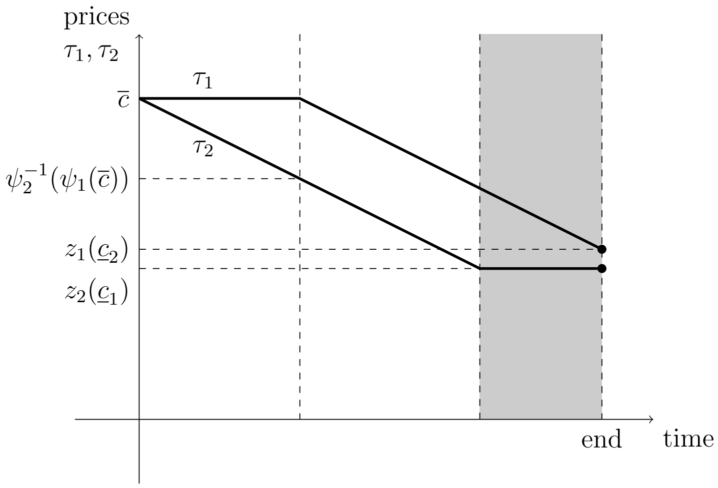

The cost vectors for which the designer greenlights project 2 despite its lower virtual surplus, represented by the shaded area in Panel 2.6a, are also represented graphically in Figure 2.7: If the auction ends in the third time segment (shaded area of Figure 2.7) before both projects can be greenlighted, project 1 must have exited because τ_1 dropped below c_1. Project 2 is greenlighted and receives transfer $z_2(\underline{c}_1)$ even though project 1 has the higher virtual surplus. Therefore if cost vectors in the shaded area of Panel 2.6a realize, the optimal descending-clock auction ends in the third time segment.

We should emphasize again a novel feature of this descending-clock auction. The clocks of both projects are paused asynchronously over some time of the auction. One project's clock runs down while the other project's clock stops. Since we have examined a very simple example, each project's clock is paused only once. In a more general setting, the projects' clocks may pause and resume several times.

Given the complexity of our problem, we do not find a simple and general ($n > 2$) full characterization of the optimal mechanism in the asymmetric case. In our examples with two projects, the problem boils down to finding one point, $(z_1(\underline{c}_2), z_2(\underline{c}_1))$, with respect to one crucial tradeoff. Naturally, the number of relevant tradeoffs increases with the number of projects. Therefore unfortunately, optimization with a larger set of projects quickly loses tractability.

Discussion

With our model as a starting point, there are several interesting modifications. In this section, we address the most natural alternative models or extensions.

$\mathbf{v}_i$ as private information, potentially correlated with $\mathbf{c}_i$ - The designer can neglect asking for v_i directly since no meaningful non-babbling equilibria in the v_i-dimension exist. If the conditional density of $v_i|c_i$ has full support, project i cannot credibly announce being a "high" type, say $\overline{v}_i$. If we slightly change the regularity assumption such that $\mathbb{E}[v_i|c_i] - c_i - \frac{F(c_i)}{f(c_i}$ must be strictly increasing, our results generalize by exchanging the previously commonly known v_i with $\mathbb{E}[v_i|c_i]$. This regularity condition mildly restricts the degree of positive correlation.

Interdependent types - We can interpret the symmetric case as a setting in which identical projects are provided at individual costs. Hence, one may wonder about a setting in which projects only draw an imperfect signal about the cost, which finally depends on other projects' signals as well. In a clock auction in such an environment, active projects update their belief about the cost whenever a project drops out. Moreover, the designer learns this information as well. Therefore the design of the optimal mechanism crucially depends on the information structure. This analysis is left for a follow-up paper.

Residual money

Whether it is reasonable to assume that the designer values residual money depends on the setting. In Ensthaler and Giebe (2014a), money does not enter the objective function, only the constraints. To clarify the relation to their paper, we introduce a linear weighting $\lambda \in [0, 1]$ of residual money, and provide comparative statics on parameter λ. The objective function can be rewritten as in (2.3),

$$\max_{\{z_i\}_{i\in I}} \mathbb{E}_{\mathbf{c}}\left[\sum_i \mathbb{I}(c_i \le z_i(\mathbf{c}_{-i}))\left(v_i - \lambda\left(c_i + \frac{F_i(c_i)}{f_i(c_i)}\right)\right)\right]$$
$$\text{s.t.}$$
$$\sum_{i\in I} \mathbb{I}(c_i \le z_i(\mathbf{c}_{-i}))z_i(\mathbf{c}_{-i}) \le B \quad \forall \mathbf{c} \in C.$$

This objective function highlights one difference to the original setting. Instead of $\boldsymbol{\zeta}^{**}$-exclusive the optimal mechanism is $\boldsymbol{\zeta}^{**}_{\lambda}$-exclusive: Define $\psi_{i,\lambda}(c) = v_i - \lambda(c + \frac{F_i(c)}{f_i(c)})$ as the λ-adjusted virtual surplus and define the vector $\boldsymbol{\zeta}^{**}_{\lambda}$ with i-the element $z^{**}_{i,\lambda} = \min\{\bar{c}_i, \psi^{-1}_{i,\lambda}(0)\}$.

It can be shown that the other properties that are sufficient to allow a DA-auction implementation continue to hold. In fact, the optimal allocation in the symmetric case remains unchanged if $\boldsymbol{\zeta}^{**}_{\lambda} = (\bar{c}_1, \bar{c}_2, \ldots, \bar{c}_n)$ for all $\lambda \in [0,1]$, i.e., when the original optimal mechanism did not exclude any cost types. For any combination of cost supports and values, there exists a sufficiently small $\lambda' > 0$ such that the designer's ranking over projects is lexicographic. In other words, λ' must be sufficiently small such that no λ'-weighted difference in cost can offset any difference in values.

FIGURE 2.8: Decreasing λ augments the quantity-quality tradeoff: The gray areas, where the project with lower λ-adjusted virtual surplus is implemented, increases.

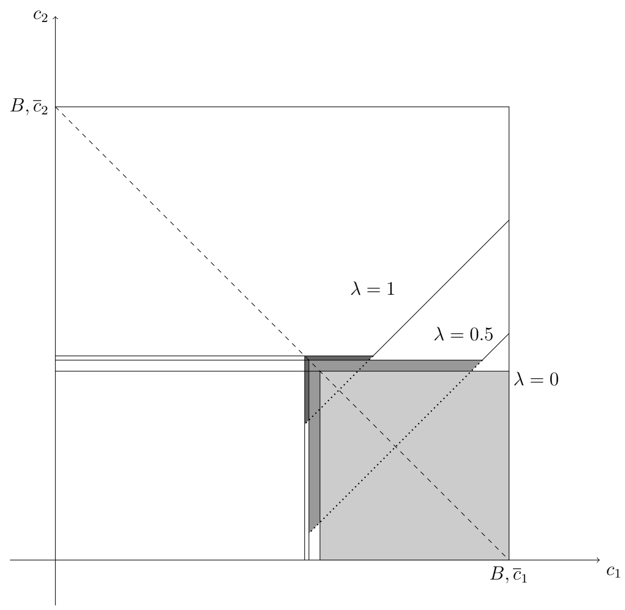

In the asymmetric case, however, the quantity-quality tradeoff is affected as well. To illustrate how the optimal allocation varies when λ is perturbed, we consider the example again, see Figure 2.8. A lower λ means that the designer prefers the high-value project 1 for higher cost reports relative to the low-value project 2 for a given cost report. This difference is illustrated by a right-shift in the diagonal that represents the loci such that both projects have equal (λ-adjusted) virtual surplus.

Reducing the weight of residual money increases the measure of cost reports for which the optimal mechanism implements project 2 despite project 1 having the larger λ-adjusted virtual surplus. Thus changing λ directly affects the quantity-quality tradeoff. As illustrated in Figure 2.8, reducing λ means that in the optimal mechanism the cutoffs at which both projects are greenlighted moves southeast, thus reducing the probability to greenlight both projects. The reason is that for lower λ a higher weight is placed on the high-value project 1.

Conclusion

Despite their importance, knapsack problems with private information have been somewhat overlooked by the economics literature. We examine a setting in which a budget-constrained procurer faces privately-informed sellers under ex-post constraints. Amongst many possible economic problems, this setting particularly applies to development funds, which are typically endowed with a fixed budget and want to finance both many projects and projects of high quality. Such problems often entail relationships in which sellers can renege on the terms of the agreement ex-post. To avoid nondelivery, shelving the project or costly renegotiation, it is appropriate to impose ex-post constraints on the agents' participation. For such settings, we have shown that a subset of DA auctions constitutes the class of optimal deterministic strategyproof mechanisms.

An optimal mechanism is described by a set of cutoff functions: All projects that report costs below their cutoff are greenlighted and receive a transfer equal to the cutoff. These cutoff functions are weakly increasing in other projects' costs, which means that the optimal allocation rule has substitutes: Given a project is implemented for some cost vector, it is also implemented when, all else being equal, the cost of a rival project is increased. Moreover, we show that the optimal allocation rule has non-bossy winners: A project that is implemented cannot affect the allocation without changing its own allocation status. In particular, if two different realizations of the cost vector lead to the same allocation, then the cutoffs of conducted projects only vary in the costs of projects not conducted. Finally, the optimal allocation rule excludes all projects with negative "virtual surplus" from the allocation.

These properties allow for a characterization as a deferred acceptance (DA) auction, introduced by Milgrom and Segal (2014). The DA auction representation provides a simple implementation via descending-clock auctions, which are easy to understand and usable in practice. In addition, DA auctions have attractive properties regarding incentive compatibility which make the prediction of equilibrium play more robust.

We fully describe the optimal allocation and the corresponding descending-clock auction in an environment in which projects are ex-ante symmetric. The optimal mechanism is monotone in the sense that the cheapest projects are greenlighted and all projects conducted receive the same transfer. This transfer either corresponds to the lowest cost among redlighted projects or the budget is distributed equally. The equivalent clock auction features a single price clock that continuously descends until all active projects can be financed.

For asymmetric environments, in which values and/or cost distributions differ, we demonstrate a novel tradeoff between quantity and quality of the greenlighted projects. The designer values both quantity and quality of the projects: She prefers projects with high virtual surplus over projects with low virtual surplus and she prefers more projects over fewer projects. In models in which the designer wants to procure a fixed number of projects, she would always choose the projects with the highest virtual surpluses. If quantity is endogenously determined by the mechanism, as in our setup, it is ex-ante not always desirable to conduct the best projects. When the best projects are always conducted, incentive compatibility would force the designer to reduce the expected number of greenlighted projects. This insight entails a consequence for the corresponding descending-clock auction. Clocks not only run asynchronously, but also periodically have to stop for certain projects.

Other interesting extensions are left for future research, for example, multiple projects per agent or projects that are complements instead of perfect substitutes. For practitioners, a simple approximately optimal mechanism may be of great value. The characterization of the optimal mechanism as a DA auction sheds light on how to construct such an approximately optimal mechanism. Halting clocks should be a key feature for the corresponding clock auction in asymmetric environments. However, we showed that the optimal strategyproof mechanism is not detail-free.

In conclusion, our methodological approach contributes to a better understanding of a class of relevant problems and opens the door for future research in this area. Furthermore, we provide an elegant indirect mechanism that can be easily implemented in practice.

Appendix

Properties of optimal mechanisms: General proofs

Lemma 2.6. *The optimal mechanism has substitutes,*

$$z_i(\widetilde{c}_j, \mathbf{c}_{-i-j}) \geq z_i(\widehat{c}_j, \mathbf{c}_{-i-j}) \text{ for almost every } \widetilde{c}_j > \widehat{c}_j \text{ and } \mathbf{c}_{-i-j} \in C_{-i-j}. \quad (2.5)$$

FIGURE 2.9: Continuous decrease / increase.

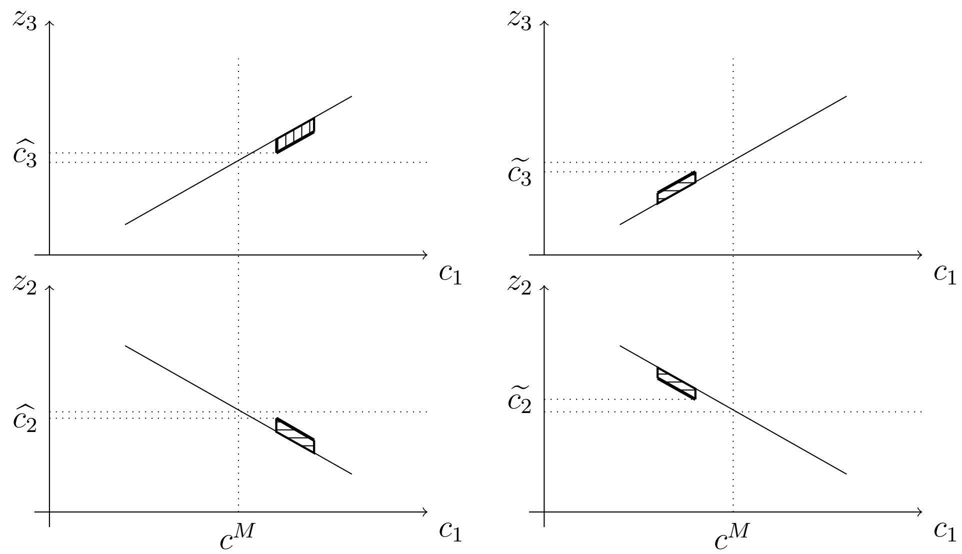

(A) Intuition for the hat deviation.

(B) Intuition for the tilde deviation.

Proof. Suppose to the contrary that somewhere z_2 is decreasing in c_1. Then there exist some ${c_1}^M$ and $\eta > 0$ such that $z_2(\underline{\underline{c_1}}, \mathbf{c}_{-1-2}) > z_2(\overline{\overline{c_1}}, \mathbf{c}_{-1-2})$ for all $\underline{\underline{c_1}} \in ({c_1}^M - \eta, {c_1}^M)$, for all $\overline{\overline{c_1}} \in ({c_1}^M, {c_1}^M + \eta)$, and for all $\mathbf{c}_{-1-2} \in \chi_{-1-2} \subset C_{-1-2}$, and χ_{-1-2} has positive Lebesgue-measure.

With more than two projects, the simple deviation of the two-project case - flattening the decreasing cutoff - is not necessarily feasible. It may be the case that other projects' cutoff functions are strictly increasing in c_1 over the same region and that for some cost vectors these cutoffs have to be paid along z_2. Then simply flattening z_2 could violate the budget constraint.

Suppose no other cutoff function is increasing while z_2 is decreasing. Then the decrease of z_2 cannot be optimal and flattening z_2 increases the designer's payoff much in the same way as in the two-project-case. Otherwise, pick a subset of

FIGURE 2.10: Jump decrease / increase.

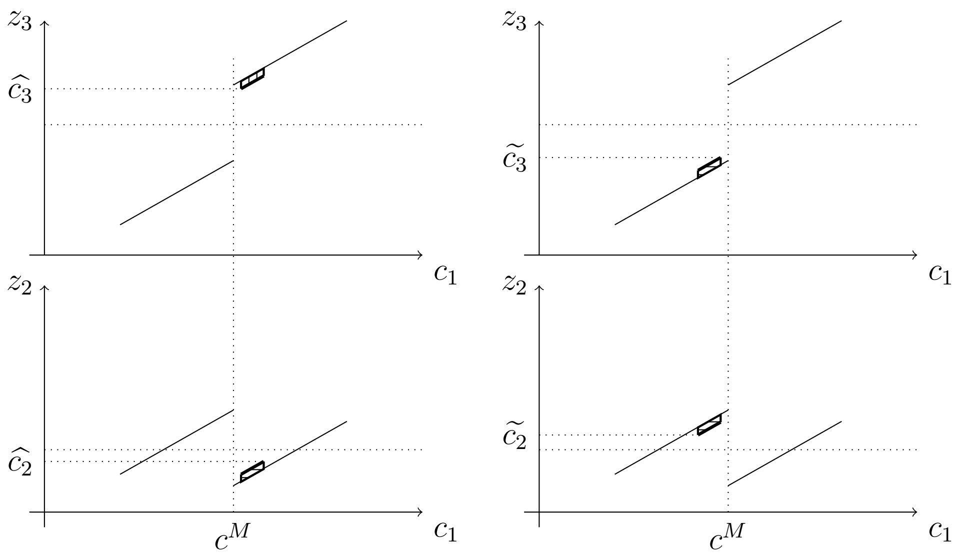

(A) Intuition for the hat deviation.

(B) Intuition for the tilde deviation.

$\widehat{\chi}_1 \subset ({c_1}^M, {c_1}^M + \eta)$ (with pos. Lebesgue-measure) such that w.l.o.g. project 3's cutoff increases in c_1 in the analogous sense to the decrease of z_2 defined above - for cost vectors where both project 2 and project 3 are eventually greenlighted, i.e., z_2 and z_3 both need to be paid.

The set

$$\begin{aligned}\widehat{\Xi}_{23}(c_1, \mathbf{c}_{-1-2-3}, \delta) = \{(c_2, c_3) | c_2 \in (z_2(c_1, c_3, \mathbf{c}_{-1-2-3}), z_2(c_1, c_3, \mathbf{c}_{-1-2-3}) + \delta]; \\ c_3 \in (z_3(c_1, c_2, \mathbf{c}_{-1-2-3}) - \delta, z_3(c_1, c_2, \mathbf{c}_{-1-2-3})]\}\end{aligned}$$

must have positive measure on $\mathbb{R}^2$ for all $c_1 \in \widehat{\chi}_1$ and for any $\mathbf{c}_{-1-2-3} \in \chi_{-1-2-3}$, where χ_{-1-2-3} is a set with positive Lebesgue measure where the cutoff of project 2 is decreasing while the cutoff of project 3 is increasing. It is the set of (c_2, c_3) tuples, where c_2 just exceeds z_2 by no more than δ, while c_3 lies just below z_3 by no more than δ - given $\mathbf{c}_{-1-2-3}$ and c_1. By $\widehat{\Xi}^2_{23}(c_1, \mathbf{c}_{-1-2-3}, \delta)$ we denote the set of project 2 components of tuples in the set $\widehat{\Xi}_{23}(c_1, \mathbf{c}_{-1-2-3}, \delta)$, and similarly for project 3.

Now deviate from the candidate mechanism in setting

$$\begin{aligned}
\widehat{z}_2(c_1, c_3, \mathbf{c}_{-1-2-3}) &:= z_2(c_1, c_3, \mathbf{c}_{-1-2-3}) + \delta \\
\widehat{z}_3(c_1, c_2, \mathbf{c}_{-1-2-3}) &:= z_3(c_1, c_2, \mathbf{c}_{-1-2-3}) - \delta \\
&\text{for all} \\
c_1 &\in (\widehat{c}_1, \widehat{c}_1 + \varepsilon) \\
c_2 &\in \widehat{\Xi}^2_{23}(c_1, \mathbf{c}_{-1-2-3}) \\
c_3 &\in \widehat{\Xi}^3_{23}(c_1, \mathbf{c}_{-1-2-3}) \\
\mathbf{c}_{-1-2-3} &\in \widehat{\chi}_{-1-2-3} \subset \chi_{-1-2-3}.
\end{aligned}$$

We call this deviation the *hat* deviation. The intuition for this deviation is the following. For an ε-environment of c_1 to the right of c_1^M (i.e., $\widehat{c}_1 > c_1^M$), increase the decreasing cutoff $z_2(c_1, c_3, \mathbf{c}_{-1-2-3})$ by δ for all c_3 that drop out of the allocation if $z_3(c_1, c_2, \mathbf{c}_{-1-2-3})$ (at c_2) is decreased by δ. Likewise only increase $z_3(c_1, c_2, \mathbf{c}_{-1-2-3})$ by δ for those c_2 that are additionally greenlighted if $z_2(c_1, c_3, \mathbf{c}_{-1-2-3})$ is increased by δ. Therefore if the deviation changes the allocation, project 2 is now greenlighted whereas project 3 is not.

This deviation is feasible. Remember that there must be enough budget to pay both z_2 and z_3 - otherwise flattening z_2 would have been possible. But then there is enough budget for $z_2 + \delta$ and $z_3 - \delta$.

Now define

$$\begin{aligned}
\widehat{c}_2 &:= \sup_{c_1, \mathbf{c}_{-1-2-3}} \widehat{\Xi}^2_{23}(c_1, \mathbf{c}_{-1-2-3}) \\
\widehat{c}_3 &:= \inf_{c_1, \mathbf{c}_{-1-2-3}} \widehat{\Xi}^3_{23}(c_1, \mathbf{c}_{-1-2-3}) \\
&\text{s.t.} \\
c_1 &\in (\widehat{c}_1, \widehat{c}_1 + \varepsilon) \\
\mathbf{c}_{-1-2-3} &\in \widehat{\chi}_{-1-2-3}.
\end{aligned}$$

In words, to bound the change in payoff we let $\widehat{c}_2$ be the highest cost type gained by the deviation and we let $\widehat{c}_3$ be the lowest cost type lost by the deviation. Then the change in payoff for the hat deviation is bounded in the following way:

$$\widehat{\Delta} > (\psi_2(\widehat{c_2}) - \psi_3(\widehat{c_3}))* \\ \int_{\widehat{\chi}_{-1-2-3}} \int_{\widehat{c_1}}^{\widehat{c_1}+\varepsilon} \int_{\widehat{\Xi}^2_{23}(c_1,\mathbf{c}_{-1-2-3})} \int_{\widehat{\Xi}^3_{23}(c_1,\mathbf{c}_{-1-2-3})} 1 dF_3(\cdot) dF_2(\cdot) dF_1(\cdot) dF_{-1-2-3}(\cdot).$$

If $\widehat{\Delta} > 0$, we have found a profitable deviation. If not, then consider the following *tilde* deviation.

Analogously to $\widehat{\Xi}_{23}$ we define the set

$$\widetilde{\Xi}_{23}(c_1, \mathbf{c}_{-1-2-3}, \delta) = \{(c_2, c_3) | c_2 \in (z_2(c_1, c_3, \mathbf{c}_{-1-2-3}) - \delta, z_2(c_1, c_3, \mathbf{c}_{-1-2-3})]; \\ c_3 \in (z_3(c_1, c_2, \mathbf{c}_{-1-2-3}), z_3(c_1, c_2, \mathbf{c}_{-1-2-3}) + \delta]\}$$

which again must have positive measure.

Now, we deviate for an ε-environment to the left of c_1^M (i.e., $\widetilde{c_1} < c_1^M$). But instead of increasing z_2 and decreasing z_3, we increase z_3 and decrease z_2:

$$\begin{aligned} \widetilde{z_2}(c_1, c_3, \mathbf{c}_{-1-2-3}) &:= z_2(c_1, c_3, \mathbf{c}_{-1-2-3}) - \delta \\ \widehat{z_3}(c_1, c_2, \mathbf{c}_{-1-2-3}) &:= z_3(c_1, c_2, \mathbf{c}_{-1-2-3}) + \delta \\ &\text{for all} \\ c_1 &\in (\widetilde{c_1} - \varepsilon, \widetilde{c_1}) \\ c_2 &\in \widetilde{\Xi}^2_{23}(c_1, \mathbf{c}_{-1-2-3}) \\ c_3 &\in \widetilde{\Xi}^3_{23}(c_1, \mathbf{c}_{-1-2-3}) \\ \mathbf{c}_{-1-2-3} &\in \widetilde{\chi}_{-1-2-3} \subset \chi_{-1-2-3}. \end{aligned}$$

The relevant bounds to bound the payoff are then given by

$$\begin{aligned} \widetilde{c_2} &:= \inf_{c_1, \mathbf{c}_{-1-2-3}} \widetilde{\Xi}^2_{23}(c_1, \mathbf{c}_{-1-2-3}) \\ \widetilde{c_3} &:= \sup_{c_1, \mathbf{c}_{-1-2-3}} \widetilde{\Xi}^3_{23}(c_1, \mathbf{c}_{-1-2-3}) \\ &\text{s.t.} \\ c_1 &\in (\widetilde{c_1} - \varepsilon, \widetilde{c_1}) \\ \mathbf{c}_{-1-2-3} &\in \widetilde{\chi}_{-1-2-3}. \end{aligned}$$

And this gives the following bound for the payoff

$$\widetilde{\Delta} > (\psi_2(\widetilde{c_3}) - \psi_3(\widetilde{c_2}))*$$
$$\int_{\chi_{-1-2-3}} \int_{\widetilde{c_1}-\varepsilon}^{\widetilde{c_1}} \int_{\widetilde{\Xi}^2_{23}(c_1,\mathbf{c}_{-1-2-3})} \int_{\widetilde{\Xi}^3_{23}(c_1,\mathbf{c}_{-1-2-3})} 1 dF_3(\cdot)dF_2(\cdot)dF_1(\cdot)dF_{-1-2-3}(\cdot).$$

By appropriately choosing δ, $\widehat{\Xi}_{-1-2-3}$, and $\widetilde{\Xi}_{-1-2-3}$, we can ensure that $\widehat{c_3} > \widetilde{c_3}$ and $\widehat{c_2} < \widetilde{c_2}$. This follows simply from the notion of increasing/decreasing cutoffs and is illustrated in Figures 2.9 and 2.10. Therefore $\widehat{\Delta} \leq 0$ implies $\widetilde{\Delta} > 0$. Consequently, there is always a profitable deviation and our candidate mechanism could not have been optimal. □

Lemma 2.8. *For any cost vectors $(\mathbf{c}_G, \mathbf{c}_R) \in C$ and $(\mathbf{c}'_G, \mathbf{c}_R) \in C$ such that $G = \gamma(\mathbf{c}_G, \mathbf{c}_R) = \gamma(\mathbf{c}'_G, \mathbf{c}_R)$ and $R = I \setminus \gamma(\mathbf{c}_G, \mathbf{c}_R)$, the optimal cutoff function z_g for all $g \in G$ is (almost everywhere) independent of the costs of all greenlighted projects $\mathbf{c}_G$. That is,*

$$z_g(\mathbf{c}_{G-g}, \mathbf{c}_R) = z_g(\mathbf{c}'_{G-g}, \mathbf{c}_R),$$

for all $\mathbf{c}_{G-g}$ and $\mathbf{c}'_{G-g}$ such that G is the set of greenlighted agents.

Proof. Take any feasible candidate mechanism with any set of increasing cutoff functions $\{z_i\}_{i\in I}$ for any individual project. Assume that for some cost vectors with positive Lebesgue-measure, only all projects in set $G \subseteq I$ are executed while all projects of set R are not conducted. Therefore there exists a set, C_R^G, with positive Lebesgue-measure containing the part of the cost vector for the projects in set R such that the partition $\{G, R\}$ is induced given some $\mathbf{c}$ where the redlighted projects have costs $\mathbf{c}_R \in C_R^G$. Then $a_i^G(\mathbf{c}_R)$ according to the following definition

$$\begin{aligned} a_i^G(\mathbf{c}_R) = \max\{c_i | \exists \mathbf{c}_{G-i} : c_i \leq z_i(\mathbf{c}_{G-i}, \mathbf{c}_R), \\ \text{and } c_g \leq z_g(\mathbf{c}_{G-j}, \mathbf{c}_{-G}) \forall g \in G, \\ \text{and } c_r > z_r(\mathbf{c}_G, \mathbf{c}_{-G-r}) \forall r \in R\} \end{aligned} \tag{2.12}$$

exists for all $i \in G$ given $\mathbf{c}_R \in C_R^G$. In words, $a_i^G(\mathbf{c}_R)$ is the highest cost of project i such that, given some cost vector $\mathbf{c}_R$ of projects that are not executed, there exists some vector $\mathbf{c}_{G-i}$ of costs of competing projects that induces a cutoff $z_i(\mathbf{c}_{G-i}, \mathbf{c}_{-G})$ above said cost while each element c_g of the vector $\mathbf{c}_{G-i}$ is lower than the cutoff induced by $a_i^G(\mathbf{c}_R)$ and the elements of the cost vectors $\mathbf{c}_R$ and $\mathbf{c}_{G-i-g}$,

$$\forall g \in G \setminus \{i\},\ \ c_g \leq z_g(\mathbf{c}_R, \mathbf{c}_{G-i-g}, a_i^G(\mathbf{c}_R)).$$

Simultaneously, it must hold that these costs induce a cutoff such that no project $r \in R$ is conducted

$$\forall r \in R,\ \ c_r > z_r(\mathbf{c}_{R-r}, \mathbf{c}_{G-i}, a_i^G(\mathbf{c}_R)).$$

Moreover, we can replace any function z_i with a left-continuous function that is identical up to a set of points with Lebesgue-measure zero. Hence, the limit is reached from below and there exists at least one cost vector $(\widehat{\mathbf{c}}_{-i}, a_i^G(\mathbf{c}_R))$ where G is the set of executed projects and $a_i^G(\widehat{\mathbf{c}}_R) = z_i(\widehat{\mathbf{c}}_{-i})$ holds. Now, notice that

$$\widehat{c}_g \leq a_g^G(\widehat{\mathbf{c}}_R)\ \forall g \in G \setminus \{i\},$$

because, given $\widehat{\mathbf{c}}_R$, there cannot exist a cost vector where only all projects in G are executed and the cost of project g exceeds $a_g^G(\widehat{\mathbf{c}}_R)$ by its construction. Moreover, we have established that every cutoff function z_i is weakly increasing in each argument. Thus,

$$a_i^G(\widehat{\mathbf{c}}_R) = z_i(\widehat{\mathbf{c}}_{-i}) \leq z_i(a_{G-i}^G(\widehat{\mathbf{c}}_R), \widehat{\mathbf{c}}_R), \tag{2.13}$$

where a_{G-i}^G is the vector of all a_g^G defined according to (2.12) except a_i^G. This inequality tells us that, whenever some vector $(\mathbf{c}_R, \mathbf{c}_{G-i}) \geq (\widehat{\mathbf{c}}_R, a_{G-i}^G(\widehat{\mathbf{c}}_R))$[15] realizes, a sufficient condition for project $i \in G$ to be executed is $c_i \leq a_i^G(\widehat{\mathbf{c}}_R)$.

The same logic also applies to all projects in G other than i. Therefore at least all projects $g \in G$ are conducted whenever a cost vector realizes such that $c_g = a_g^G(\mathbf{c}_R)$.[16] Consequently, the budget constraint requires that

$$\sum_{g \in G} z_g(a_{-g}^G(\mathbf{c}_R), \mathbf{c}_R) \leq B. \tag{2.14}$$

Furthermore, given $\mathbf{c}_R$, for all projects $g \in G$, $z_g(\mathbf{c}_{-G}, \mathbf{c}_R) = a_g^G(\mathbf{c}_R)$ if $\mathbf{c}_{G-g} \leq a_{G-g}^G(\mathbf{c}_{-G})$. That is, the cutoffs are constant given the cost vector of redlighted projects.

[15] When $\mathbf{x}$ and $\mathbf{y}$ are vectors, $\mathbf{x} \geq \mathbf{y}$ means that every element x_i of $\mathbf{x}$ weakly exceeds the corresponding element y_i of $\mathbf{y}$.

[16] $a_i^G(\mathbf{c}_R)$ is only defined if $C^G \neq \emptyset$ and $\mathbf{c}_R \in C_R^G$, but this does not hinder the proof.

Suppose to the contrary that $z_i(\mathbf{c}_{-i}) < a_i(\mathbf{c}_R)$ for some $i \in G$ and for all $\mathbf{c}_{-i} \in \Xi \subset C^G_{-i}$ with Ξ having positive Lebesgue measure.

Define $\Xi(\mathbf{c}_{G-i-j}, \mathbf{c}_R) \subset [0, \bar{c}_j]$ where $z_i(\mathbf{c}_{G-i-j}, c_j, \mathbf{c}_R) < a^G_i(\mathbf{c}_R)$ for all $c_j \in \Xi(\mathbf{c}_{G-i-j}, \mathbf{c}_R)$. For any $\mathbf{c}_{G-i-j} \leq a^G_{-i-j}(\mathbf{c}_R)$, let

$$z^{\Xi}_i(\mathbf{c}_{G-i-j}, \mathbf{c}_R) := \max_{c_j \in \Xi(\mathbf{c}_{G-i-j}, \mathbf{c}_R)} z_i(\mathbf{c}_{G-i-j}, c_j, \mathbf{c}_R)$$

By (2.14), changing the mechanism to

$$z_i(\mathbf{c}_{G-i,-j}, c_j, \mathbf{c}_R) = a^G_i(\mathbf{c}_R), \quad \forall c_j \leq a^G_j(\mathbf{c}_R)$$

does not violate the budget constraint. This deviation increases the payoff conditional on $\mathbf{c}_R$ by

$$\Delta > \int_{\Xi_{-j}} \Pr(c_j \in \Xi(\mathbf{c}_{G-i-j}, \mathbf{c}_R)) \int_{z^{\Xi}_i(\mathbf{c}_{G-i-j}, \mathbf{c}_R)}^{a^G_i(\mathbf{c}_R)} \psi_i(c) dF_i(c) dF_{-i-j}(\mathbf{c}_{-i-j}) > 0.$$

Given that Ξ has positive Lebesgue-measure, this deviation also strictly increases the unconditional payoff. □

Constructing a scoring function: Proof of Proposition 2.11

To prove Proposition 2.11, it is helpful to consider the following lemmata. While Lemma 2.8 (non-bossy winners) is a statement that conditions on a fixed allocation, it also has implications on the cutoffs resulting from different cost vectors that induce different allocations.

Lemma 2.16. *Take any mechanism and any two cost vectors $\mathbf{c} \neq \widehat{\mathbf{c}}$ that induce partitions $\{G, R\}$ and $\{\widehat{G}, \widehat{R}\}$, respectively. Then*

$$\mathbf{c}_{R \cup \widehat{R}} = \widehat{\mathbf{c}}_{R \cup \widehat{R}}$$
$$\mathbf{c}_{G \cap \widehat{G}} \neq \widehat{\mathbf{c}}_{G \cap \widehat{G}}$$

implies

$$G = \widehat{G}$$
$$R = \widehat{R},$$

that is, $\mathbf{c}$ *and* $\widehat{\mathbf{c}}$ *induce the same allocation.*

Proof. Given cost vector $\mathbf{c}$, define a new cost vector $\mathbf{c}'$, where $c_i' = \min\{c_i, \widehat{c}_i\}$ for all $i \in G \cap \widehat{G}$ and $\mathbf{c}'_{R \cup \widehat{R}} = \mathbf{c}_{R \cup \widehat{R}}$. By Lemma 2.8, $\mathbf{c}'$ induces allocation $\{G, R\}$. Similarly, a perturbation of cost vector $\widehat{\mathbf{c}}$ in the same way with $\widehat{c}_i' = \min\{c_i, \widehat{c}_i\}$ for all $j \in G \cap \widehat{G}$ and $\widehat{\mathbf{c}}'_{R \cup \widehat{R}} = \widehat{\mathbf{c}}_{R \cup \widehat{R}}$ must induce allocation $\{\widehat{G}, \widehat{R}\}$. But $\mathbf{c}' = \widehat{\mathbf{c}}'$ by construction. Hence, $G = \widehat{G}$ and $R = \widehat{R}$. □

Lemma 2.17. *Take any mechanism and any two cost vectors* $\mathbf{c} \neq \widetilde{\mathbf{c}}$ *that induce partitions* $\{G, R\}$ *and* $\{\widetilde{G}, \widetilde{R}\}$*, respectively. Then*

$$z_i(\mathbf{c}_{G \cap \widetilde{G}}, \mathbf{c}_{R \cup \widetilde{R}}) = z_i(\widetilde{\mathbf{c}}_{G \cap \widetilde{G}}, \mathbf{c}_{R \cup \widetilde{R}})$$
$$z_j(\widetilde{\mathbf{c}}_{G \cap \widetilde{G}}, \widetilde{\mathbf{c}}_{R \cup \widetilde{R}}) = z_j(\mathbf{c}_{G \cap \widetilde{G}}, \widetilde{\mathbf{c}}_{R \cup \widetilde{R}})$$

for all $i \in G$ *and for all* $j \in \widetilde{G}$*, respectively.*

Proof. By Lemma 2.16, the vector $(\widetilde{\mathbf{c}}_{G \cap \widetilde{G}}, \mathbf{c}_{R \cup \widetilde{R}})$ leads to allocation $\{G, R\}$ and the vector $(\mathbf{c}_{G \cap \widetilde{G}}, \widetilde{\mathbf{c}}_{R \cup \widetilde{R}})$ leads to allocation $\{\widetilde{G}, \widetilde{R}\}$. The rest follows directly from Lemma 2.8 (non-bossy winners). □

Having established these properties we can prove Proposition 2.11 by induction. We construct a DA scoring function for each iteration. Conditional on all previous iterations having been constructed correctly, we can demonstrate how to construct an appropriate scoring function for any iteration.

Proposition 2.11. *Any optimal mechanism has a DA auction representation and can be implemented with a descending-clock auction.*

Proof. This proof is structured as follows. First, we construct scoring functions for each iteration of the DA auction. Then we explain how the zeros of the scoring functions are derived. Finally we show by induction that the constructed DA auction implements the same allocation as the underlying z-mechanism.

Scoring functions

First, we introduce some notation. Let A_t be the set of active projects in iteration t and let $O_t := I \setminus A_t$ be the set of inactive projects (O as in "out"). Let $O_t j := O_t \cup \{j\}$ be the union of dropped out projects and some individual project j.

Fix an optimal z-mechanism and consider the corresponding DA auction with scoring functions $\{s_i^A\}_{A\subset I, i\in A}$

$$s_i^A(c_i, \mathbf{c}_O) = \begin{cases} 0 & \text{if } c_i \leq a_i^A(\mathbf{c}_O), \\ c_i + \sum\limits_{i\neq j}^{j\in A} b_{Oi}^j(c_i, \mathbf{c}_O) & \text{otherwise,} \end{cases} \tag{2.15}$$

where $a_i^A(\mathbf{c}_O)$ is defined as in (2.12) and $b_{Oi}^j(c_i, \mathbf{c}_O)$ is defined as

$$\begin{aligned} b_{Oi}^j(c_i, \mathbf{c}_O) :=& \max\left\{c_j : \exists \tilde{\mathbf{c}}_{-Oi-j} : R = Oi | c_i, \mathbf{c}_O\right\} \\ :=& \max\Big\{c_j : \exists \tilde{\mathbf{c}}_{-Oi-j} : c_i > z_i(c_j, \tilde{\mathbf{c}}_{-Oi-j}, \mathbf{c}_O), \\ & \quad \text{and } c_o > z_o(c_i, c_j, \tilde{\mathbf{c}}_{-Oi-j}, \mathbf{c}_{O-i}) \forall o \in O, \\ & \quad \text{and } c_g \leq z_g(c_i, c_j, \tilde{\mathbf{c}}_{-Oi-j}, \mathbf{c}_O) \forall g \in A \setminus i\Big\}. \end{aligned}$$

In words, $b_{Oi}^j(c_i, \mathbf{c}_O)$ is the highest cost of project j such that given the vector $\mathbf{c}_{Oi}$ the corresponding z-mechanism implements the allocation partition $R = Oi$ and $G = A \setminus i$ for some realization of the cost vector $\tilde{\mathbf{c}}_{-Oi-j}$.

Zeros of the scoring functions

Suppose the DA auction ends in the t-th iteration. Then all projects $i \in A_t$ have score $s_i^{A_t} = 0$ and the cost vector must induce $G = A_t$ in the underlying z-mechanism. By non-bossiness of winners, cutoffs of projects in G are constant in the part of the cost vector $\mathbf{c}_{A_t}$ for all cost vectors inducing the same allocation.

Therefore we can characterize the zeros of the scoring function by a threshold and $s_i^{A_t} = 0$ whenever project i's cost is below this threshold. The threshold is given by $a_i^{A_t}(\mathbf{c}_O)$ as defined in (2.12). Notice that $c_i \leq a_i^{A_t}(\mathbf{c}_O)$ implies that project i is not eliminated in the t-th iteration, even if other projects exceed their threshold. This implication does not rule out permissible z-mechanisms. Conditional on $\mathbf{c}_{O_t}$, some projects exceeding their threshold can at most lead to a higher cutoff for project i due to monotonicity.

Further notice that if $c_i > a_i^{A_t}(\mathbf{c}_O)$, there always exist cost vectors with $\mathbf{c}_{O_t}$ for previously eliminated projects that induce $G = A_t \setminus \{i\}$. For example, all cost vectors with $c_j \leq a_j^{A_t}(\mathbf{c}_{O_t})$ for all $j \in A_t \setminus \{i\}$ induce that allocation. However, this condition is sufficient for $G = A_t \setminus \{i\}$ but not necessary. There can be other cost vectors inducing the same allocation.

Iteration 1

If multiple projects have a positive score, it also holds that

$$\text{If } \widehat{\mathbf{c}} \text{ induces } \widehat{R} = \{i\} \text{ then } s_i^I(c_i) > s_j^I(c_j) \text{ for all } j \neq i \tag{2.16}$$

The meaning of $\widehat{R} = \{i\}$ is that $\widehat{c}_i > z_i(\widehat{\mathbf{c}}_{-i})$ and $\widehat{c}_j \leq z_j(\widehat{\mathbf{c}}_{-j})$. Hence, by construction

$$\widehat{c}_j \leq z_j(\widehat{\mathbf{c}}_{-j}) \leq b_i^j(\widehat{c}_i) \tag{2.17}$$

as $b_i^j(\widehat{c}_i)$ is the highest cutoff z_j that allows allocation $\widehat{R} = \{i\}$ given $\widehat{c}_i$.

Next, we show

$$\widehat{c}_i > b_j^i(\widehat{c}_j). \tag{2.18}$$

Suppose that the contrary holds, then there exists a vector $\widetilde{\mathbf{c}}_{-i-j}$ such that

$$\widehat{c}_i \leq z_i(\widehat{c}_j, \widetilde{\mathbf{c}}_{-i-j})$$

and allocation $\widetilde{R} = \{j\}$ is implemented. By Lemma 2.17 we know that the cutoffs z are constant in costs of projects $\widehat{G} \cap \widetilde{G} = I \setminus \{i, j\}$. Consequently, we arrive at

$$\widehat{c}_i \leq z_i(\widehat{c}_j, \tilde{\mathbf{c}}_{-i-j}) = z_i(\widehat{c}_j, \widehat{\mathbf{c}}_{-i-j})$$

which means that i is greenlighted for vector $\widehat{\mathbf{c}}$, a contradiction to our initial assumption that $\widehat{\mathbf{c}}$ implements $\widehat{R} = \{i\}$.

Next, we show

$$b_i^k(\widehat{c}_i) \geq b_j^k(\widehat{c}_j) \text{ for all } j \neq i \text{ and } k \neq i, j. \tag{2.19}$$

By definition

$$b_i^k(\widehat{c}_i) = z_k(\widehat{c}_i, \tilde{\mathbf{c}}_{-i-k}) \text{ for some } \tilde{\mathbf{c}}_{-i-k},$$
$$b_j^k(\widehat{c}_j) = z_k(\widehat{c}_j, \dot{\mathbf{c}}_{-j-k}) \text{ for some } \dot{\mathbf{c}}_{-j-k}.$$

Because projects $-i-j-k$ are greenlighted for both cost realizations $(\widehat{c}_k, \widehat{c}_i, \tilde{\mathbf{c}}_{-i-k})$ and $(\widehat{c}_k, \widehat{c}_i, \dot{\mathbf{c}}_{-i-k})$, it follows by Lemma 2.17 that

$$b_i^k(\widehat{c}_i) = z_k(\widehat{c}_i, \tilde{\mathbf{c}}_{-i-k}) = z_k(\widehat{c}_i, \widehat{\mathbf{c}}_{-i-k}),$$
$$b_j^k(\widehat{c}_j) = z_k(\widehat{c}_j, \dot{\mathbf{c}}_{-j-k}) = z_k(\dot{c}_i, \widehat{\mathbf{c}}_{-i-k}).$$

Furthermore, it must hold that $\widehat{c}_i > \dot{c}_i$, otherwise vector $\widehat{\mathbf{c}}$ would not optimally redlight project i while vector $(\widehat{\mathbf{c}}_{-\mathbf{i}}, \dot{c}_i)$ optimally greenlights project i. Then by bidder substitutability,

$$b_i^k(\widehat{c}_i) = z_k(\widehat{c}_{-k}) \geq z_k(\dot{c}_i, \widehat{\mathbf{c}}_{-i-k}) = b_j^k(\widehat{c}_j).$$

Combining (2.17), (2.18) and (2.19) leads to (2.16). We have shown that the scoring function eliminates the correct project when $|R| = 1$, i.e., the redlighted project.

Finally, we need to show that if $|R| > 1$, the project removed in the first iteration is redlighted in the allocation implemented by the underlying z-mechanism, i.e.,

$$A_1 \setminus A_2 = \{k\} \Rightarrow k \in R.$$

Now take cost vector $\widetilde{\mathbf{c}}$ with allocation $\{\widetilde{G}, \widetilde{R}\}$ and let $i \in \widetilde{G}$ be some greenlighted project and and let $j \in \widetilde{R}$ be some redlighted project, respectively. Since project j is redlighted, it must have cost $\widetilde{c}_j > a_j^I$. Hence there exists some cost vector $\widehat{\mathbf{c}}$ with $\widehat{c}_j = \widetilde{c}_j$ such that $\widehat{R} = \{j\}$. By Lemma 2.17, we can assume $\widehat{c}_i = \widetilde{c}_i$ since $i \in \widetilde{G} \cap \widehat{G}$. As our scoring function correctly matches all cases in which $|R| = 1$, it must be that $s_j(\widetilde{c}_j) > s_i(\widetilde{c}_i)$. Given that we have chosen i and j arbitrarily, we have shown that any project removed in the first iteration must be in the redlighted set, which was to show.

Iteration 2

We can show with the same arguments as above, that the previously stated scoring function is correct for $t = 2$ as well. To this end, we inductively rely on the fact that the project k removed in the first iteration is indeed redlighted by the z-mechanism - as we have shown above.

Iteration $t \geq 3$

With the appropriate scoring functions used in all previous iterations, we can then show that the t-th iteration removes the correct project for all cost vectors inducing $|R| = t$ given a z-mechanism and otherwise removes some project $i \in A_t$, where $i \in R$, for all cost vectors inducing $|R| > t$. □

The symmetric case

Proposition 2.12. *Arrange the projects in ascending order of their reported costs, $c_1 \leq c_2 \leq \cdots \leq c_n \leq c_{n+1} := \bar{c}$, and define $z^k := \min\left\{\frac{B}{k}, z^{**}, c_{k+1}\right\}$. In the symmetric case, the cutoff mechanism with $z_i(\mathbf{c}_{-i}) = z^{k^*}$ is the optimal mechanism. The optimal number of accepted projects k^* is given by $k^* := \max\{k | c_k \leq z^k\}$.*

Proof. The case $n = 2$ has been proven in Section III.i.

Now, consider $n = 3$. Fix any c_3 and any mechanism as candidate for optimality. Either $c_3 > z_3(c_1, c_2)$ or $c_3 \leq z_3(c_1, c_2)$. In the first case, project 3 is not executed and the budget remaining for the other two is still B. In the second case, project 3 is executed and the budget remaining for the other two becomes $B - z_3(c_1, c_2)$.

Now, consider deviating to the proposed mechanism only for project 1 and 2. The change in profit looks like a probability weighted sum of terms similar to the two-project case, only that the distributions F are conditional on c_1 and c_2 being in some interval (that induces $z_3 >$ or $< c_3$) and the budget must be adjusted.

Because log-concavity of F implies log-concavity of $\frac{F(c)-F(a)}{F(b)-F(a)}$ this deviation is always positive like in the case $n = 2$. The same logic can be applied to any n, changing any mechanism by selecting two projects and then adjusting their cutoffs in the following way: The budget is shared equally if both projects are executed; if only one project is executed, it has to be the one with higher virtual surplus; never execute projects with negative virtual surplus. Iterating over these steps ultimately arrives at the proposed mechanism which has to be optimal. □

Bidder Substitutability and Complementarity

In the main text, we made the crucial assumption that $\underline{c} = 0$. As a consequence, complementaries as in the following example are excluded. The example shows that an optimal mechanism may not have substitutes. When the lower bound of all projects' costs is zero, it is always possible to improve a mechanism that does not have substitutes. The idea of the proof of optimality of bidder substitutability in Appendix A is that it cannot be optimal to decrease i's cutoff to the benefit of increasing j's cutoff when some third project's cost increases from c_k to $c_k' > c_k$, because then it would either be better to raise project $z_j(\cdot, c_k)$ at the cost of lowering $z_i(\cdot, c_k)$ as well or it would be better to raise $z_i(\cdot, c_k')$ at the cost of lowering $z_j(\cdot, c_k')$.

In the following example, this approach is not feasible. Through the lower cost bounds and the values, projects 1 and 2 inherit endogenous complementarities. The designer prefers implementing 1 and 2 together over implementing 3 alone, but once either 1 or 2 becomes too expensive the other project is dropped as well in favor of implementing only project 3.

Example 2.2. *Suppose $I = \{1, 2, 3\}$ and $B = 300$. Let the costs be arbitrarily distributed on the following supports:*

$$c_1 \sim [200, 400], c_2 \sim [20, 200], c_3 \sim [290, 300],$$

and let the values be

$$v_1 = 700, v_2 = 500, v_3 = 1000.$$

Let the corresponding optimal mechanism be given by

$$z_1(c_2, c_3) = \begin{cases} 250 \text{ if } c_2 \leq 50 \\ 0 \text{ otherwise} \end{cases}, z_2(c_1, c_3) = \begin{cases} 50 \text{ if } c_1 \leq 250 \\ 0 \text{ otherwise} \end{cases},$$

$$z_3(c_1, c_2) = \begin{cases} 300 \text{ if } c_2 > 50 \text{ or } c_1 > 250 \\ 0 \text{ otherwise} \end{cases}.$$

Bidder substitutability fails because, e.g., as c_1 increases from 249 to 251, project 2 with, say, cost 40 gets dropped from the allocation set. The designer cannot, as in the main text with $\underline{c} = 0$, lower $z_3(40, 249)$ as it is already zero or profitably raise $z_2(251, \cdot)$ at the cost of project 3 as the lower cost bounds prohibit that projects 2 and 3 are ever conducted together and implementing $G = \{3\}$ is preferred to $G' = \{2\}$.

However, it is still possible to construct an implementation with price clocks: All clocks start at the upper bounds. Then (at arbitrary speed) the prices of 1 and 2 decrease to $(250, 50)$. If both projects are still active, the price for project 3 decreases to zero while clocks 1 and 2 halt: 1 and 2 are implemented. If any project $i \in \{1,2\}$ drops out earlier, then the price for $j \neq i$, $j \in \{1,2\}$ drops to zero, while price 3 remains at 300. 3 is implemented.

The next example features another kind of complementatity. In this example project 3 can be a bossy loser. Again, there exists a DA-auction implementation. The lower cost bounds of the (stochastically) identical projects 1 and 2 are too high for both projects to ever be conducted together. The cheaper of the two is greenlighted. Project 3 is then only implemented if enough money remains.

Example 2.3. *Suppose $I = \{1,2,3\}$ and $B = 300$. Let the costs be arbitrarily distributed on the following supports:*

$$c_1, c_2 \sim [151, 200], c_3 \sim [50, 300],$$

and let the values be

$$v_1 = v_2 = 1000, v_3 = 500.$$

Let the corresponding optimal mechanism be given by

$$z_1(c_2, c_3) = c_2, \quad z_2(c_1, c_3) = c_1, \quad z_3(c_1, c_2) = B - \max\{c_1, c_2\}.$$

Suppose $c_2 > c_1$, then project 2 can be a bossy loser: It can increase its cost report without changing its status to the green light and thereby kick project 3 out of the allocation.

While substitutes and non-bossiness are sufficient for an implementation with a DA auction, they are clearly not necessary. From the matching literature, it is apparent that some kind of substitutes condition is needed and non-bossy winners seem to be important for DA implementations. We have constructed a scoring function that implements the exemplary allocations above. However, in the proof of non-bossiness of winners, we need the strong substitutes condition for inequality (2.13).

A weaker substitutes condition, such as our groupwise substitutes, does not suffice for the optimality of non-bossy winners. This condition is satisfied by the examples above and is helpful for the construction of a scoring function.

Definition 2.18. An allocation rule γ has groupwise substitutes, if $\sum_{g \in G} z_g(c_{-g})$ is increasing in any cost report c_r with $r \notin G$ for all allocation sets G that are admitted by γ.

3. Groups, cheap talk, and voting

Introduction

In settings in which several individuals have to make a collective decision, voting is a very common if not the most common decision-making process. As such, voting is a method to aggregate private information from different individuals. In some cases, however, in particular when voters act strategically, voting might not be the best way to aggregate information. Feddersen and Pesendorfer (1998), for example, point out this inferiority of voting in the context of a jury in a criminal trial. When voting is the only way to aggregate private information about guilt or innocence and unanimity is needed to convict, the resulting outcome can be an excess of convictions. However, in their model this problem vanishes once jurors are allowed to talk prior to voting. Given that talking to each other is usually possible, this example serves to demonstrate the importance of the strategic interaction of information aggregation through voting and through talking, i.e., through cheap talk.

Often committees that vote on an issue consist of different groups. Political parties in a parliamentary committee are one example of such a situation. On such a committee, the way in which individuals talk to each other is different from the way they talk to each other in the absence of groups. Normally, members of the same party would first talk amongst their fellow party members before talking to members of other parties.

To investigate the interaction between this form of group-to-group communication and voting, I consider a model of a reform decision made by a committee consisting of two groups, a majority with 3 members and a minority with 2 members. Each individual receives a (binary) private signal about the (binary) state of the world, which is whether the reform would be a success if implemented. A successful reform gives every individual the same benefit but implementing the reform is associated with different costs for the majority and the minority. Costs and benefit of the reform are such that all individuals would want to implement a successful reform and not implement an unsuccessful reform, if the state of the world was known. Both groups are risk-neutral and weigh the probability of success of the reform against their costs of implementing the reform.

For my analysis, I consider each group to act as a player in a two-player game. The implicit assumption in this model choice is that both groups pool their private information within the group and then perfectly coordinate their strategies afterwards. As both groups are perfectly homogeneous, this assumption is plausible and the resulting equilibrium also exists in a game in which each individual is a single player instead of the group. However, less plausible equilibria in which groups fail to coordinate are ruled out.

The timing is as follows. First, both groups can simultaneously send a cheap talk message about their groups' signals to the other group. Then both groups simultaneously vote, where the number of votes of each group equals its number of members. Depending on the voting rule, the reform either needs a majority or a unanimous pro-reform vote to pass.

For this game, I characterize the set of implementable deterministic social choice functions, i.e., the mapping from signals to the outcome, and specify strategies with which any such social choice function can be implemented. These strategies are such that only one group conveys information through cheap talk. Thus no social choice function requires for implementation that both groups talk instead of babbling. Moreover, if the set of implementable social choice functions is not a singleton, one social choice function must Pareto-dominate the others. This result allows me to compare outcomes under the two different voting rules: majority rule and unanimity. Interestingly, there exist parameters such that both groups strictly prefer majority rule and other parameters such that both groups strictly prefer unanimity. Consequently, understanding whether a committee consists of different groups and, if so, how these groups communicate can have important implications for the choice of the voting procedure.

There is an extensive literature on the influence of cheap talk (Crawford and Sobel, 1982) prior to voting. The seminal paper by Austen-Smith (1990) is one of the earliest contributions in this strand of the literature, of which a good overview is provided by Austen-Smith and Feddersen (2009). The role of voting thresholds with pre-voting communication has been analyzed by Gerardi and Yariv (2007). They very elegantly point out that the set of implementable outcomes is not affected by the choice of the voting rule, as long as the rule is non-unanimous. However, their argument hinges on achieving an outcome using voting strategies with which no single voter is ever pivotal. Thus all voting strategies are necessarily equilibrium strategies. This result, however, is no longer plausible when there exist voter groups that can potentially coordinate their behavior.

With majority rule, my model is close to the model of Ishida and Shimizu (2016), who consider cheap talk between an informed sender and an informed receiver, where the latter makes a decision affecting both. Their setup is therefore directly related to settings in my model in which the minority sends a cheap talk message to the majority who can then decide afterwards. However, in contrast to me, they focus on conditions for full communication, in particular when sender and receiver receive signals with different precision.

The importance of considering groups in the context of communication and voting is also not completely novel. In their overview, Austen-Smith and Feddersen (2009) call for research in this area. To my knowledge, Thordal-Le Quement and Yokeeswaran (2015) and Hummel (2012) have provided the only studies with this focus, apart from this paper. Both Hummel (2012) and Thordal-Le Quement and Yokeeswaran (2015) consider within-group cheap talk followed by voting. Contrary to Thordal-Le Quement and Yokeeswaran (2015) and me, Hummel (2012) focuses on the large committee limit. He shows that in the limit, the presence of groups who can communicate via within-group cheap talk can lead to the probability of making a correct decision going to 1.

Probably most closely related to this paper is the recent work of Thordal-Le Quement and Yokeeswaran (2015), who also consider a committee with two groups. The crucial difference is that they consider communication within a group followed by coordinated voting by both groups while I consider group-to-group cheap talk prior to coordinated voting, which is a richer communication protocol. For any equilibrium under their communication protocol there exists one under my communication protocol giving the same outcome. However, the converse is not true. Keeping the voting rule - unanimity - fixed, they then make a welfare comparison between three different communication protocols: subgroup deliberation (cheap talk within the group), plenary deliberation (cheap talk among all committee members), and no deliberation. They find that the outcome with the first protocol Pareto-dominates the outcome with the second protocol, which Pareto-dominates the outcome with the third protocol. I, on the other hand, keep the communication protocol fixed (group-to-group cheap talk) and make a welfare comparison between two voting rules: majority rule and unanimity.

Finally, the information structure in my model can be considered the work horse of the cheap talk and voting literature. Albeit usually framed a bit differently in the context of a jury trial, the same structure of binary private signals, a binary state of the world, and a binary collective decision is used by Feddersen and Pesendorfer (1998), Doraszelski, Gerardi, and Squintani (2003), Ishida and Shimizu (2016) (when we interpret sender and receiver as groups), Thordal-Le Quement and Yokeeswaran (2015), Hummel (2012), and many others. Thus, the basic structure of their models and my model is either the same or at least isomorphic for relevant parameter values.

Model

A committee consisting of two groups has to make a binary decision $a \in \{0, 1\}$. This decision is about a reform with $a = 1$ signifying the reform's implementation. From an ex-ante perspective, the reform, if implemented, will be successful, $s = 1$, with probability $\frac{1}{2}$. There are two groups, majority M with 3 group members and minority m with 2 group members. I number the individuals starting with the majority. An individual i will thus be from the majority group if $i \in \{1, 2, 3\}$ and i will be from the minority group if $i \in \{4, 5\}$. The assumption about the size of groups M and m is made for the sake of exposition. I conjecture that one can also derive the main results of this paper with larger groups, albeit with vastly bigger case distinctions in some of the proofs.

The benefit from the proposed reform will depend on the true state of the world $s \in \{0, 1\}$ and is identical for all individuals from both groups. The benefit is normalized to 1 if the reform is successful ($s = 1$) and 0 if the reform fails ($s = 0$). However, individuals from the majority and individuals from the minority have to bear different costs for implementing the reform. If the reform is implemented, denoted by $a = 1$, members of the majority bear cost $c_M \in [0, 1]$ and members of the minority bear cost $c_m \in [0, 1]$ - regardless of the reform's success. Hence, the ex-post utility function of group $x \in \{M, m\}$ is given by

$$u_x = a(s - c_x).$$

Each individual i is endowed with a private signal $s_i \in \{0, 1\}$ that is correlated with the true state of the world with precision p, i.e. $Pr(s_i = 1|s = 1) = Pr(s_i = 0|s = 0) = p$, $p \in (0.5, 1)$. As all signals are drawn with the same precision, an individual from any group weighs all signals he can observe equally and the relevant statistic to judge the reform is the sum of positive signals. To an individual who observes k out of n positive signals the posterior probability of the reform's success and thus the expected benefit from the reform to this individual is given by:

$$\beta(k, n) := \frac{p^k(1-p)^{n-k}}{p^k(1-p)^{n-k} + (1-p)^k p^{n-k}}.$$

As individuals in both groups are risk-neutral, an individual that could observe all 5 signals wants to see the reform implemented if

$$\beta(k, 5) - c_x \geq 0,$$

where $x \in \{M, m\}$ denotes the individual's group and $k = \sum_{i=1}^{5} s_i$.

The eventual decision whether to implement the reform is made by voting. Each individual can cast a vote $a_i \in \{0, 1\}$ and the reform is implemented if the number of affirmative votes reaches threshold τ:

$$a = 1 \Leftrightarrow \sum_i a_i \geq \tau.$$

I use two different voting rules, *majority rule* with $\tau = 3$ and *unanimity* with $\tau = 5$.

Voting occurs after an initial round of communication. The goal of this paper is to illuminate the interaction of information transmission between different groups with different incentives through cheap talk and through subsequent voting. To this end, from here on, I consider each group to act as a single economic agent. Thus both M and m are agents who only differ in the number of individuals they represent and therefore in the number of signals and in the number of votes they are endowed with. For readability, I use female pronouns for M and male pronouns for m. Here,

$$s_M := \sum_{i=1}^{3} s_i \in \{0,1,2,3\} = \mathcal{S}_M, \qquad s_m := \sum_{i=4}^{5} s_i \in \{0,1,2\} = \mathcal{S}_m,$$

denote the number of positive signals observed by M and m, respectively. Thus I call s_M M's signal and s_m m's signal. Alternatively, one can interpret s_M and s_m as M's and m's respective types in a common values public good decision setup.

While at first the restriction of considering both groups as two agents seems perhaps parsimonious, it does not really exclude any interesting equilibria because the interests within each group are perfectly aligned. Therefore, it is not unreasonable to assume that group members within each group act in unison. Consequently, all the results of the following analysis can also be derived in an alternative model in which instead each individual is an economic agent, and individuals in both groups can first communicate within their group and then from group to group. In the relevant equilibria both groups pool their information, and then coordinate their behavior. This restriction does, however, rule out some implausible equilibria which hinge on the groups failing to coordinate and often can only be supported by a construction in which no single individual is pivotal.

With respect to both cheap talk and voting, the timing is as follows. First, both groups observe their respective signals. Then both groups simultaneously exchange messages about s_M and s_m. The majority announces message $r_M \in \mathcal{S}_M$ and the minority announces message $r_m \in \mathcal{S}_m$. Then the both groups simultaneously cast their affirmative votes, $a_M \in \mathcal{S}_M$ for the majority and $a_m \in \mathcal{S}_m$ for the minority. The reform is implemented if $a_M + a_m \geq \tau$, in which case the success of the reform s is revealed and payoffs realize.

Analysis

In this section, I derive all my results with respect to the model described in the previous section. Statements that require a proof (Lemma 1, Lemma 2, Lemma 3, and Proposition 3.2) are proved in the Appendix. Other statements either follow directly from previously established results (Corollary 3.1) or are shown by example (Proposition 3.3). The relevant equilibrium concept for this analysis is Perfect Bayesian equilibrium. To obviate the need for equilibrium refinements I

FIGURE 3.1: An example with the same threshold of doubt $q_M = q_m = 3$ without a conflict of interest (top) and with different thresholds of doubt $q_M = 2 < q_m = 3$ and thus with a conflict of interest (bottom). A circle with number x marks the posterior $\beta(x, 5)$ if x out of 5 positive signals realize for precision $p = 0.66$.

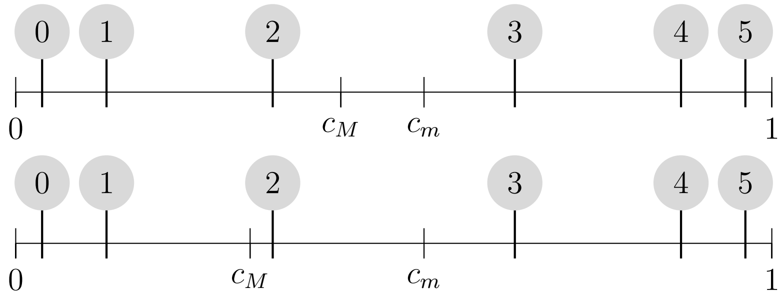

only consider equilibria in which all available messages are played on path in the cheap talk stage. With cheap talk, this restriction is without loss of generality, as for any message strategy that does not use all available messages there exists an alternative strategy that uses all messages and conveys the same amount of information.[1]

A useful concept that is often employed in models with the same type of information structure as in this model (e.g., Feddersen and Pesendorfer, 1998; Thordal-Le Quement and Yokeeswaran, 2015) is the notion of the threshold of reasonable doubt q_x. It is defined as the lowest number of positive signals out of the total number of available signals such that group x wants to implement the reform. Therefore, a threshold $q_x > 0$ satisfies

$$\beta(q_x - 1, 5) < c_x \leq \beta(q_x, 5).$$

Information sharing in the context of this model means making the outcome contingent on the signals of both groups. Depending on the mode of voting, τ, information can be shared between both groups using communication, using voting, or using both. Consider an example with majority voting, $\tau = 3$. Suppose in this example m wants to allow M to condition the outcome on his own signal s_m. In particular, M wants to implement the reform whenever $s_m \geq 1$. One way to achieve this outcome is to reveal s_m at the communication stage, i.e., $r_m = s_m$. Then M could vote $a_M = 3$ whenever $r_m \geq 1$, achieving the desired outcome. Alternatively, m could simply cast a number of affirmative votes corresponding

[1]For a more detailed version of this argument see Farrell (1993, Section 3).

to his signal, i.e., $a_m = s_m$. Then, $a_M = 2$ also achieves the desired outcome, implementing the reform whenever $s_m \geq 1$, regardless of the information conveyed at the communication stage.

In terms of payoff in equilibrium, it is not important whether information was conveyed through cheap talk or through voting. As information sharing through cheap talk and through voting are at least to a degree substitutable, there can be several equilibria with different strategies at the cheap talk and at the voting stage that are equivalent in terms of payoff. Instead, the equilibrium can uniquely be characterized payoff-wise by a social choice function α, mapping from both group's signals into outcomes. For the sake of tractability, I only consider equilibria that implement a deterministic social choice,

$$\alpha : \{0,1,2,3\} \times \{0,1,2\} \to \{0,1\}.$$

This restriction is not without loss of generality. In particular, it rules out equilibria in which there is mixing at the voting stage between two different strategies that could lead to different outcomes. However, I believe that the restriction is comparatively innocent in this setting as there could only be a limited amount of mixing in any case.

To illustrate this point, consider the following example in a setup with majority rule ($\tau = 3$). Suppose that m plays a mixed strategy at the voting stage for some realization of s_m, e.g., $s_m = 1$, and after receiving some message r_M. Suppose m mixes between $a_m = 2$ and $a_m = 1$. Mixing only meaningfully affects the outcome if M sometimes votes $a_M = 1$ given s_M and r_m, so that $a_m = 2$ and $a_m = 1$ lead to different outcomes. Therefore, m must be indifferent whether the reform is implemented given his information, $s_m = 1, r_M$, and what he can infer about s_M from M voting $a_M = 1$. But if M is indifferent for $s_m = 1$ he must strictly prefer the reform not to be implemented for $s_m = 0$ and he must strictly prefer the reform to be implemented if $s_m = 2$. Consequently, conditional on the information transmitted before the voting stage, m can only mix between $a_m = 1$ and $a_m = 2$ for one of the three possible realizations of s_m in a way that affects the outcome.

Additionally, the decision to limit attention to deterministic social choice functions has the advantage of preserving comparability to Thordal-Le Quement and Yokeeswaran (2015). For those equilibria with the communication protocol that is closest to my setup, subgroup deliberation (SD) equilibria, they consider two groups that first pool their information within the group and then all vote for reform (or in their case for conviction) if the number of positive signals is weakly greater than some threshold and all vote against otherwise. By construction, this class of equilibria leads to a deterministic social choice function. Moreover, the set of social choice functions that can be implemented by SD equilibria is a subset of the set of implementable social choice functions. The additional cheap talk between groups with my communication protocol does not eliminate SD equilibria, as both groups can simply babble and then follow the equilibrium voting strategy of an SD equilibrium. Therefore, all their welfare comparisons between SD equilibria and other equilibria under other communication protocols also hold in my model, whenever their and my model coincide.[2]

As established above, information sharing can occur both through cheap talk and through voting. Therefore, the amount of information shared in equilibrium can only be meaningfully quantified based on the social choice function α. To this end, I distinguish between *information transmission*, *full information transmission* and *partial information transmission*.[3]

Definition 1 (Information transmission). An equilibrium has **information transmission** if it implements a social choice α such that

$$\begin{aligned} \alpha(s_M, s_m) &\neq \alpha(s'_M, s_m) \text{ for some } s_M, s'_M, s_m \\ \alpha(s_M, s_m) &\neq \alpha(s_M, s'_m) \text{ for some } s_m, s'_m, s_M. \end{aligned}$$

Having an equilibrium with information transmission means that the social choice function is contingent on both groups' signals. Thus information from both groups is used in determining the outcome.

Definition 2 (Full information transmission). An equilibrium satisfies **full information transmission** if there exists an equilibrium implementing the same social choice α in which both groups truthfully reveal their signals at the communication stage in equilibrium.

[2]The models do not always coincide, as Thordal-Le Quement and Yokeeswaran (2015) assume $\max\{c_M, c_m\} > 0.5$. In other words, based solely on the prior, the group paying a greater cost for reform dislikes the reform. In terms of their court application, the doves on the jury favor acquittal without information. They do not claim that this assumption is necessary, but they use it in their proofs.

[3]Thordal-Le Quement and Yokeeswaran (2015) have similar categories. Their notion of a *fully reactive* equilibrium coincides with my notion of an equilibrium for which the corresponding social choice function has (partial) information transmission.

In contrast, full information transmission is stronger than information transmission.[4] An equilibrium has full information transmission if no group would want to deviate at the voting stage, even if it was fully informed about the other group's signal. Therefore, each group either receives or can infer all the information it seeks from the other group. It is well known (e.g., Coughlan, 2000, Proposition 5) that full information transmission can generally only be achieved if M and m have the same preferences, i.e., $q_M = q_m$.[5] While the full information transmission case is straight-forward, the interesting case for this analysis is the case of partial information transmission, the case in which information transmission is feasible but full information transmission is not.

Definition 3 (Partial information transmission). An equilibrium satisfies **partial information transmission** if the social choice α exhibits information transmission but not full information transmission.

In order to examine partial information transmission social choice functions, I first need to establish if and when such social choice functions can be implemented in equilibrium. This is not a trivial task, since I have shown that there can be different combinations of communication and voting strategies that implement the same social choice function α. Therefore, showing that some set of strategies that would implement α is not an equilibrium does not automatically rule out other strategies that might implement α in equilibrium.

This is not a Mechanism Design problem, as the rules of the game that both groups play are fixed. Nevertheless, I follow an approach similar to the revelation principle approach in Mechanism Design to tackle this problem. First, I formulate necessary conditions that the social choice function α must satisfy in any equilibrium. Then I characterize strategies that implement α and show that these strategies constitute an equilibrium whenever the necessary conditions are fulfilled. Therefore, these conditions are both necessary and sufficient for α to be implementable in equilibrium.

[4] I assume $q_M, q_m \in \{1, 2, 3, 4, 5\}$ in this statement. This assumption rules out situations in which group x is never interested in the other group's signal because it is always for the reform ($q_x = 0$) or always against the reform ($q_x > 5$). If the other group cannot be swayed, then it is an equilibrium strategy to truthfully reveal one's signal. Thus such a social function would have the silly property of having *full information transmission* but not *information transmission*, in the sense of the definitions.

[5] Full information transmission can also be achieved in the non-generic case in which $|q_M - q_m| = 1$ and one of the groups, group x, is indifferent with respect to the reform's implementation if q_x positive signals realize: $c_x = \beta(q_x, 5)$.

The first condition any α must satisfy to be implementable in equilibrium is monotonicity (MON). If the reform is implemented given signals (s_M, s_m), it must also be implemented for any signals (s'_M, s'_m) that are weakly larger in both dimensions.

$$\begin{aligned}
&\alpha(s'_M, s_m) \geq \alpha(s_M, s_m) \qquad \forall s'_M, s_M \in \mathcal{S}_M; s'_M > s_M; s_m \in \mathcal{S}_m \\
&\alpha(s_M, s'_m) \geq \alpha(s_M, s_m) \qquad \forall s'_m, s_m \in \mathcal{S}_m; s'_m > s_m; s_M \in \mathcal{S}_m
\end{aligned} \tag{MON}$$

An equivalent statement is that any monotone α can be represented by threshold rules $\psi_M(s_m)$ and $\psi_m(s_M)$ for each group. Then, for example, the reform is implemented given s_m only when $s_M \geq \psi_M(s_m)$.

$$\begin{aligned}
\psi_M(s_m) &:= \min\{\widehat{s}_M | \alpha(\widehat{s}_M, s_m) = 1\} \\
\psi_m(s_M) &:= \min\{\widehat{s}_m | \alpha(s_M, \widehat{s}_m) = 1\}
\end{aligned}$$

The necessity of monotonicity follows almost directly from the fact that the reform decision is monotone in the number of yes-votes. To see this point, suppose monotonicity fails. Then, receiving a larger signal must induce one group to cast a lower number of yes-votes, after having received some message of the other group. Suppose that, after having received some message r_m, M votes a_M given s_M but votes $a'_M < a_M$ given $s'_M > s_M$. Voting a_M leads to reform whenever a'_M leads to reform. In addition, for monotonicity to fail a_M must lead to reform in some cases in which a'_M does not. As M votes a_M with signal s_M, the posterior belief in the reform's success must exceed c_M, given s_M, given m's message, and given what M can infer about s_m when a_M leads to reform but a'_M does not, Thus M wants to chose a_M instead of a'_M. But she also wants to do so given $s'_M > s_M$, contradicting our initial assumption that she votes a_M.

Apart from (MON), any equilibrium must satisfy the incentive constraints (IC_M) and (IC_m). These constraints ensure that M and m weakly prefer following their equilibrium strategies given signal s_M or s_m rather than deviating and following the strategy prescribed for some other signal $\widehat{s}_M$ or $\widehat{s}_m$.

$$\sum_{s_m \in \mathcal{S}_m} \Pr(s_m|s_M)\alpha(s_M, s_m)(\beta(s_m + s_M, 5) - c_M) \geq \tag{IC_M}$$
$$\sum_{s_m \in \mathcal{S}_m} \Pr(s_m|s_M)\alpha(\widehat{s}_M, s_m)(\beta(s_m + s_M, 5) - c_M) \qquad \forall s_M, \widehat{s}_M \in \mathcal{S}_M$$
$$\sum_{s_M \in \mathcal{S}_M} \Pr(s_M|s_m)\alpha(s_M, s_m)(\beta(s_m + s_M, 5) - c_m) \geq \tag{IC_m}$$
$$\sum_{s_M \in \mathcal{S}_M} \Pr(s_M|s_m)\alpha(s_M, \widehat{s}_m)(\beta(s_m + s_M, 5) - c_m) \qquad \forall s_m, \widehat{s}_m \in \mathcal{S}_m$$

The final set of constraints depends on the choice of τ. With majority rule, $\tau = 3$, the constraints (IR^1_M) and (IR^0_M) ensure that the majority does not want to deviate given s_M and unilaterally prevent the reform - ensured by (IR^1_M) - or unilaterally force the reform - ensured by (IR^0_M) - regardless of s_m. For unanimity, $\tau = 5$, I need constraints (IR^1_M) and (IR^1_m), as both groups can prevent the reform unilaterally but neither can force it:

Additional constraints with majority rule:

$$\sum_{s_m \in \mathcal{S}_m} \Pr(s_m|s_M)\alpha(s_M, s_m)(\beta(s_m + s_M, 5) - c_M) \geq 0 \qquad \forall s_M \in \mathcal{S}_M \tag{IR^1_M}$$

$$\sum_{s_m \in \mathcal{S}_m} \Pr(s_m|s_M)(1 - \alpha(s_M, s_m))(\beta(s_m + s_M, 5) - c_M) \leq 0 \qquad \forall s_M \in \mathcal{S}_M \tag{IR^0_M}$$

Additional constraints with unanimity:

$$\sum_{s_m \in \mathcal{S}_m} \Pr(s_m|s_M)\alpha(s_M, s_m)(\beta(s_m + s_M, 5) - c_M) \geq 0 \qquad \forall s_M \in \mathcal{S}_M \tag{IR^1_M}$$

$$\sum_{s_M \in \mathcal{S}_M} \Pr(s_M|s_m)\alpha(s_M, s_m)(\beta(s_m + s_M, 5) - c_m) \geq 0 \qquad \forall s_m \in \mathcal{S}_m \tag{IR^1_m}$$

Having established necessary conditions for α to be implementable in some equilibrium, I can construct specific strategies that implement any α in equilibrium, as long as α satisfies these conditions. These strategies are stated in Lemma 1.

Lemma 1. With the following strategies any α that is feasible given majority rule

($\tau = 3$) and any α that is feasible given unanimity ($\tau = 5$) can be implemented in equilibrium of the communication and voting game with the appropriate τ:

1. M babbles at the communication stage by mixing over all messages with equal probability.

2. Given s_m, m mixes over all messages $r_m \in \{\widehat{s}_m | \alpha(s_M, \widehat{s}_m) = \alpha(s_M, s_m) \forall s_M \in \mathcal{S}_M\}$ with equal probability.

3. M votes $a_M = 3$ if $\alpha(s_M, r_m) = 1$ and $a_M = 0$ if $\alpha(s_M, r_m) = 0$.

4. m votes $a_m = 2$ at the voting stage unless $\alpha(s_m, s_M) = 0$ for all $s_M \in \mathcal{S}_M$ given s_m.

Lemma 1 is useful insofar as it allows me to now only consider the strategies from the lemma for the remaining analysis. Moreover, in any equilibrium constructed with the strategies described in the lemma, m uses one out of four distinct communication strategies. These strategies are described in Corollary 3.1: A strategy for full information transmission, two distinct strategies for partial information transmission, termed *communication strategy (a)* and *communication strategy (b)*, and a strategy for babbling. Consequently, for the focus of this analysis, equilibria with partial information transmission, it is sufficient to consider two different strategies, which considerably simplifies the analysis. Figure 3.2 illustrates the construction of such an equilibrium under majority rule.

Corollary 3.1. *Any social choice α can be implemented in an equilibrium in which the minority m uses one of the following four communication strategies:*

1. ***Full information transmission:*** $r_m = s_m$

2. ***Communication strategy (a):***

$$r_m = \begin{cases} 0 & \textit{if } s_m = 0 \\ \in \{1, 2\} \textit{ with equal prob.} & \textit{if } s_m \in \{1, 2\} \end{cases}$$

3. ***Communication strategy (b):***

$$r_m = \begin{cases} \in \{0, 1\} \textit{ with equal prob.} & \textit{if } s_m \in \{0, 1\} \\ 2 & \textit{if } s_m = 2 \end{cases}$$

4. ***Babbling:*** *$r_m \in \{0, 1, 2\}$ with equal probability*

In a first step, I can use Lemma 1 and Corollary 3.1 to investigate the limits of information transmission. Like full information transmission, partial information transmission is limited by the degree of conflict between M and m. While full information transmission is only feasible when preferences are perfectly aligned, i.e. when $q_M = q_m$, partial information transmission is only feasible if preferences are sufficiently aligned. Lemma 2 gives the necessary condition for partial information transmission in this setup, $|q_M - q_m| \leq 1$. This condition is not general, but specific to the setup in this paper. With larger groups, this condition, i.e., the maximal difference of q_M and q_m, will increase, as a difference of the same size represents a smaller degree of conflict with larger groups.

Lemma 2. Information transmission can only occur in equilibrium as long as $|q_M - q_m| \leq 1$.

Having established Lemma 2, I derive Lemma 3 in order to further characterize how an implementable social choice α is implemented with the strategies described in Lemma 1. This result is useful as it allows me to focus on a manageable number of possible equilibria in the proof of Proposition 3.2.

Lemma 3. In an equilibrium that implements a social choice α with partial information transmission and which is constructed with the strategies described in Lemma 1, the communication strategy of m depends on τ, q_M, and q_m in the following way.

With majority rule ($\tau = 3$):

- If $q_m < q_M$: m uses communication strategy (a).
- If $q_m > q_M$: m uses communication strategy (b).

With unanimity ($\tau = 5$):

- If $q_m < q_M$: m uses communication strategy (a).
- If $q_m > q_M$: Either communication strategy can occur.

Proposition 3.2 is a crucial result that makes all subsequent observations possible. Both models with voting and and with cheap talk generally allow for a multitude of equilibria. As a result, there might be more than one social choice function that is implementable in equilibrium. Therefore, I need to select one social choice function to make meaningful comparisons between outcomes for different cost parameters or to compare voting with majority rule and with unanimity.

Proposition 3.2. *If several social choice functions are implementable in equilibrium, there is always a Pareto-dominant one.*

Proposition 3.2 establishes that if several social choice functions are implementable for a given setup, one always Pareto-dominates the others. This social choice is thus a natural candidate for selection. In all subsequent comparisons I only consider Pareto-dominant implementable social choice functions. Thordal-Le Quement and Yokeeswaran (2015, Lemma 6 and Proposition 8) have a comparable argument in their setup. While conceptually similar, it does not, however, easily extend to my setup.

With unanimity the set of implementable social choice functions with my group-to-group communication protocol rule is a superset of the implementable social choice functions with their subgroup deliberation (SD) communication protocol. Therefore, the best implementable outcome in my setup weakly Pareto-dominates the best outcome in their setup. Now I can use their main results (Propositions 7-9), which state that the outcome with the SD communication protocol Pareto-dominates the outcome with a plenary deliberation (PD) protocol, where cheap talk messages are received by all individuals (not just their own group). This outcome in turn Pareto-dominates the outcome with a no deliberation (ND) protocol, where no cheap talk is permitted. Therefore, whenever $\max\{c_M, c_m\} > 0.5$ (one of their assumptions), the outcome under my group-to-group protocol not only Pareto-dominates the SD outcome, but also the PD and ND outcomes.

The proof of Proposition 3.2 consists of an extensive and thus slightly tedious case distinction, in order to show that whenever several social choice functions are implementable, one of those social choices Pareto-dominates the others from an ex-ante perspective. In all cases, similar arguments can be used to prove this statement, building on a combination of some of the relevant constraints - (IC_M), (IC_m), (IR_M^1), (IR_M^0), and (IR_m^1). Unfortunately, even though the proofs in all these cases are similar, there does not seem to be a way to easily generalize the argument in order to construct a more elegant proof.

For the conclusion of the analysis, all subsequent observations are made using graphs. Here, I use the parameter $p = 0.66$. However, one can produce qualitatively equivalent graphs with other values of p such that $\frac{1}{2} < p < 1$.

My first set of observations is about the extent of information transmission. Figure 3.3 shows the parameters such that information transmission social choice functions are feasible both for majority rule and unanimity. The dark square areas represent (c_M, c_m)-values such that $q_M = q_m$ and M and m can therefore fully share information. The outlined white rectangles represent (c_M, c_m)-values such that $|q_M - q_m| = 1$. As established in Lemma 2, information transmission is not feasible for (c_M, c_m)-values that don't lie in either area. However, while $|q_M - q_m| \leq 1$ is a necessary condition for the Pareto-dominant social choice function to have information transmission, it is not a sufficient condition. The gray areas in Figure 3.3 represent (c_M, c_m)-values with partial information transmission and clearly only cover part of the areas with $|q_M - q_m| = 1$.

My second set of observations is about partial information transmission being more prevalent with unanimity when compared to majority rule. In Figure 3.3, this fact is reflected by the gray area being larger in the right panel when compared to the left panel. This observation implies that if c_M and c_m are drawn independently from a uniform distribution, information transmission is more likely to occur with unanimity than with majority rule. Figure 3.4 also depicts this result. The dark areas in its left panel show (c_M, c_m)-values for which there is information transmission with majority rule but not with unanimity, and vice-versa in its right panel. While the overall area is larger for unanimity, there are also dark areas for majority rule. Therefore, it is not the case that there is information transmission with unanimity, whenever there is information transmission with majority rule.

Figure 3.5 illustrates how unanimity in general benefits the minority. The dark areas in the left panel depict (c_M, c_m)-values for which the majority strictly benefits from majority rule when compared to unanimity. The dark areas in the right panel depict (c_M, c_m)-values for which the minority strictly benefits from unanimity rule when compared to majority rule. Both areas mostly overlap, implying that in most cases in which the majority is worse off with unanimity compared to majority rule the minority is better off, and vice-versa.

However, both areas do not overlap fully. Therefore, surprisingly, there are cases in which both groups are better off with either majority rule or unanimity. The (c_M, c_m)-values for both cases are depicted in Figure 3.6. Thus, depending on the parameters, the majority can potentially benefit from relinquishing decision power and the minority from relinquishing veto power. For emphasis, I state this fact in Proposition 3.3.

Proposition 3.3. *There exist parameters c_M, c_m, and p, such that both M and m strictly prefer majority rule over unanimity in terms of the ex-ante payoffs of the Pareto-dominant social choice functions. Conversely, there exist parameters such that both M and m prefer unanimity over majority rule.*

A comparison of Figures 3.4 and 3.6 reveals that the result of Proposition 3.3 is not or at least not entirely due to majority rule or unanimity allowing for information transmission, where the other voting rule does not allow for it, i.e., it is not due to making the outcome contingent on both groups' outcome. Graphically, the dark areas in Figure 3.6 do not lie within the dark areas in Figure 3.4. In some cases, both majority rule and unanimity allow for information transmission, while, nevertheless, the outcome with one of the voting rules strictly dominates the outcome with the other voting rule for both groups.

Moreover, whenever majority rule or unanimity allows for information transmission, while the other voting rule does not, the outcome under the voting rule with information transmission does not necessarily Pareto-dominate the outcome under the voting rule without. Graphically, the dark areas in Figure 3.4 do not lie within the dark areas in Figure 3.6.

Finally, one can directly observe that the outcome with majority rule and the outcome with unanimity only differ when the minority is less pro-reform than the majority, i.e., if $q_m > q_M$. Here, giving the minority m the power to veto the reform does not change the outcome. Clearly, whenever M wants the reform, m wants the reform as well and therefore does not veto. Consequently, switching from majority rule to unanimity does not affect the outcome in those cases. In Figures 3.3 - 3.6, this fact results in both the left panel and the right panel being equal below an imaginary 45°-line, or $(c_M = c_m)$-line. Below this line $c_M > c_m$ and thus $q_M \geq q_m$. However, if unanimity was needed for keeping the status quo instead of achieving the reform, this picture would be reversed, i.e., both panels would be equal above the $(c_M = c_m)$-line.

Conclusion

In this paper, I contribute to the literature that investigates the interplay of cheap talk and voting in the context of a common values setup. While this field in general has been studied extensively, there is little research on how the presence of different groups affects these models, when each group consists of individuals with the same preferences (think about political parties, hawks vs. doves, constitutional originalists vs. constitutional evolutionists in US courts, etc.). To the best of my knowledge, Thordal-Le Quement and Yokeeswaran (2015), Hummel (2012), and I are the only ones that have this particular focus so far. My focus, in particular, lies on the communication between homogeneous groups - that have previously pooled their information - and how this communication interacts with their voting strategies.

In my model, I consider two homogeneous groups who first communicate and then vote on whether to pass a reform requiring either a majority or unanimity in favor. Given potential equilibrium multiplicity, I analyze the model in terms of implementable social choice functions. In particular, I derive conditions for the implementability of *information transmission* social choice functions, social choice functions with which the outcome is contingent on both groups' outcome. Not surprisingly, information transmission is only feasible when preferences are sufficiently close.

Moreover, I find that there always exists a Pareto-dominant and thus focal social choice function. This result allows me to compare the outcomes under two different voting rules: unanimity and majority rule. For most parameters, unanimity benefits the minority and majority rule benefits the majority. However, there exist parameters for which majority rule strictly Pareto-dominates unanimity and vice-versa. Therefore, the majority can benefit from relinquishing decision power or the minority can benefit from relinquishing veto power.

FIGURE 3.2: An example of partial information exchange under majority rule: $c_M = 0.7$, $c_m = 0.5$, $p = 0.66$. The top row illustrates the degree of conflict of interest. The second row contains m's and M's strategies. The third and fourth rows illustrate the incentive compatibility of M's voting strategy and m's reporting strategy, respectively.

$q_M \neq q_m \quad \Rightarrow \quad$ Full information exchange is not feasible.

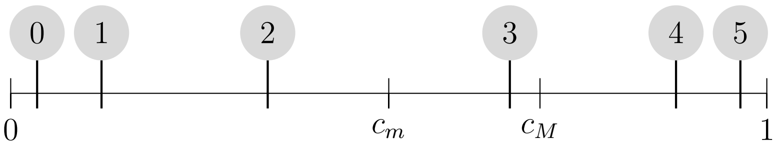

m reports:

$$s_m = 0 \Rightarrow r_m = 0$$
$$s_m \in \{1, 2\} \Rightarrow r_m \in \{1, 2\}$$

m votes (without influencing the outcome):

$$a_m = 2$$

M babbles:

$$s_M \in \{0, 1, 2, 3\} \Rightarrow r_M \in \{0, 1, 2, 3\}$$

M votes:

$$r_m = 0 \vee s_M < 2 \Rightarrow a_m = 0$$
$$r_m \in \{1, 2\} \wedge s_M \geq 2 \Rightarrow a_m = 3$$

M's voting strategy is incentive compatible given m's reporting strategy. The message $r_m = 0$ (light shaded circles) or $r_m \in \{1, 2\}$ (dark shaded circles) shifts M's posterior. Combined with M's own signal, s_M, this leads to his voting strategy:

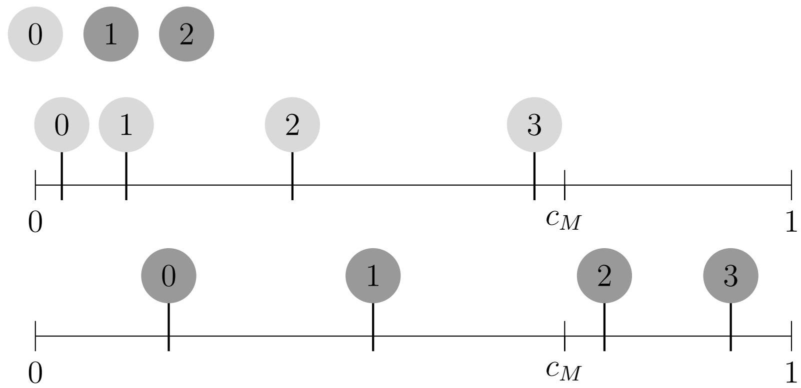

m's communication strategy is incentive compatible given M's voting strategy. m's report is pivotal if $s_M \geq 2$ in which case m wants M to vote pro-reform if $s_m > 0$:

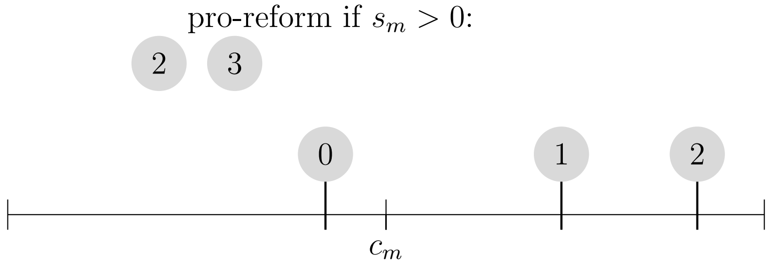

FIGURE 3.3: Degrees of information transmission under majority rule and unanimity (with $p = 0.66$). The dark squares along the diagonal represent parameters for which $q_M = q_m$, allowing for full information transmission. The outlined white rectangles represent parameters for which $|q_M - q_m| = 1$ and partial information transmission is potentially feasible. The gray rectangles represent parameters for which social choice functions with partial information transmission are actually implementable.

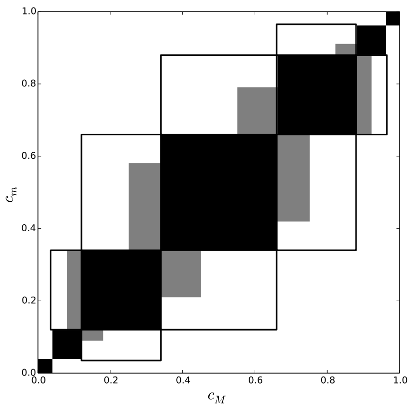

(A) Majority rule

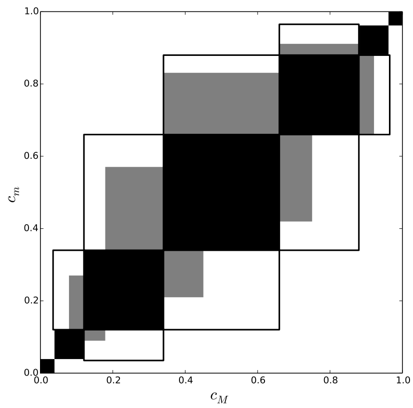

(B) Unanimity

FIGURE 3.4: Comparing information transmission with majority rule and with unanimity (with $p = 0.66$). The dark areas in the left panel represent parameters for which the social choice with majority rule has information transmission but the social choice with unanimity does not. Conversely, the dark areas in the left panel represent parameters for which the social choice with unanimity has information transmission but the social choice with majority rule does not.

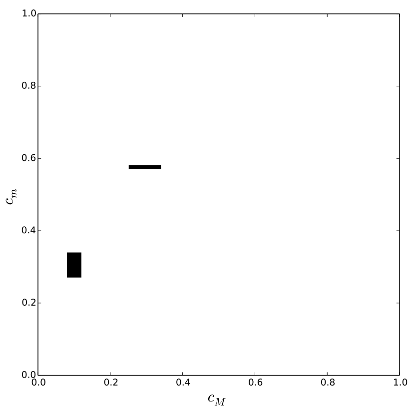

(A) Majority rule

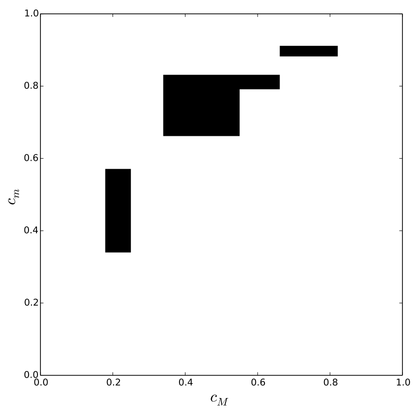

(B) Unanimity

FIGURE 3.5: Comparing payoffs with majority rule and with unanimity (with $p = 0.66$). The dark area in the left panel represents parameters for which M's payoff is higher with majority rule compared to unanimity. Conversely, the dark in the right panel represents parameters for which m's payoff is higher with unanimity compared to majority rule.

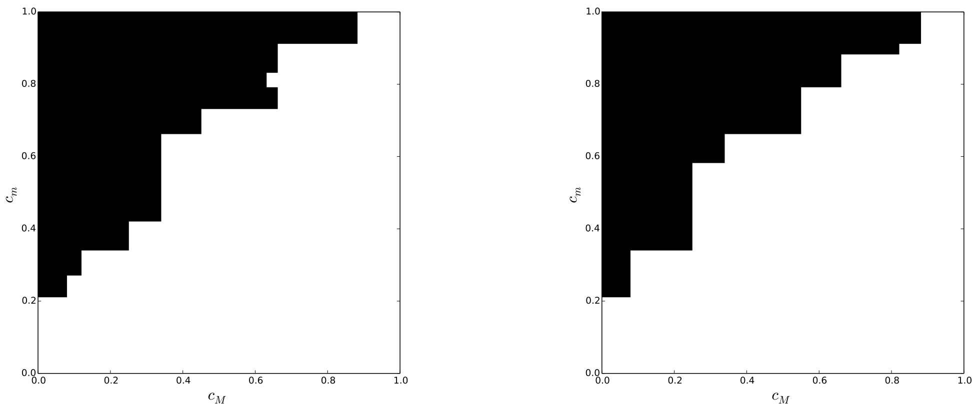

(A) M's payoff higher with majority rule

(B) m's payoff higher with unanimity

FIGURE 3.6: Comparing payoffs with majority rule and with unanimity (with $p = 0.66$). The dark areas in the left panel represent parameters for which the outcome under majority rule strictly Pareto-dominates the outcome under unanimity. Conversely, the dark areas in the right panel represent parameters for which the outcome under unanimity strictly Pareto-dominates the outcome under majority rule.

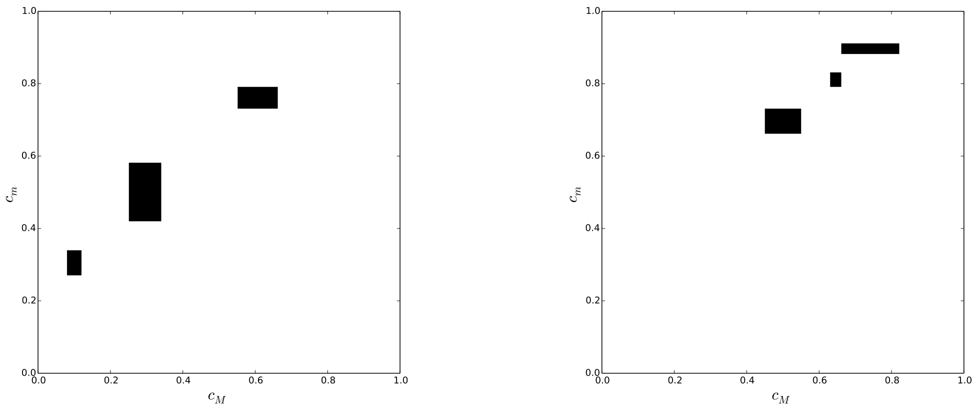

(A) Majority rule

(B) Unanimity

Appendix

Lemma 1. With the following strategies any α that is feasible given majority rule ($\tau = 3$) and any α that is feasible given unanimity ($\tau = 5$) can be implemented in equilibrium of the communication and voting game with the appropriate τ:

1. M babbles at the communication stage by mixing over all messages with equal probability.
2. Given s_m, m mixes over all messages $r_m \in \{\widehat{s}_m | \alpha(s_M, \widehat{s}_m) = \alpha(s_M, s_m) \forall s_M \in \mathcal{S}_M\}$ with equal probability.
3. M votes $a_M = 3$ if $\alpha(s_M, r_m) = 1$ and $a_M = 0$ if $\alpha(s_M, r_m) = 0$.
4. m votes $a_m = 2$ at the voting stage unless $\alpha(s_m, s_M) = 0$ for all $s_M \in \mathcal{S}_M$ given s_m.

Proof. It is easily verified that the stated strategies lead to the outcome prescribed by α. Thus, to prove the lemma I need to show that (IC_M), (IC_m), (IR^1_M), and (IR^0_M) (with majority rule) and (IC_M), (IC_m), (IR^1_M), and (IR^0_M) (with unanimity) are not only necessary conditions but also sufficient to guarantee that the stated strategies are equilibrium strategies.

First, clearly M has no incentive to deviate at the communication stage since deviating does not affect the outcome. With majority rule, m has no incentive to deviate at the voting stage, as deviating again does not affect the outcome. With unanimity, m deviating at the voting stage affects the outcome whenever the equilibrium strategies require the reform to be implemented. Then, m's expected utility from voting yes conditional on s_m and conditional on m's vote being pivotal is given by

$$\frac{\sum_{s_M \in \mathcal{S}_M} \Pr(s_M|s_m)\alpha(s_M, s_m)(\beta(s_m + s_M, 5) - c_m)}{\sum_{s_M \in \mathcal{S}_M} \Pr(s_M|s_m)\alpha(s_M, s_m)} \geq 0$$

which is greater than zero, the expected utility of voting no, due to (IR^1_m). Therefore, what remains to be shown is that m does not have a profitable deviation at the communication stage and that M does not have a profitable deviation at the voting stage.

First, suppose to the contrary of what is to be shown that m, given s_m, has a profitable deviation $\widehat{r}_m$. Then there must be some $s_M \in \mathcal{S}_M$ such that $\alpha(s_M, \widehat{r}_m) \neq \alpha(s_M, s_m)$. Thus, given M's voting strategy, m's interim payoff is given by

$$\sum_{s_M \in \mathcal{S}_M} \Pr(s_M|s_m)\alpha(s_M, \widehat{r}_m)(\beta(s_m + s_M, 5) - c_m).$$

But by (IC_m), this payoff must be weakly lower than m's payoff from following his equilibrium strategy. Therefore, the deviation cannot be profitable.

Second, consider M's strategy at the voting stage. After having received message r_m, M can infer that s_m must lie in the set $X = \{s_m|\alpha(s_M, s_m) = \alpha(s_M, r_m)\forall s_M \in \mathcal{S}_M\}$.

Suppose M always votes no regardless of r_m. Then, m babbles by construction. With majority rule, the left-hand side of (IR^0_M) represents M's expected payoff from deviating and voting yes, which is smaller than zero, M's expected payoff from voting no. Thus, (IR^0_M) ensures that M's strategy is optimal. With unanimity, whenever M always votes no m also always votes no by construction. Therefore, deviating and voting yes does not affect the outcome. Now suppose M always votes yes regardless of r_m. Then again m babbles and the left-hand side of (IR^1_M) represents M's expected payoff when always voting yes. So (IR^1_M) ensures that M's strategy is optimal.

For the final case, suppose M sometimes votes yes and sometimes votes no depending on s_M and r_m. Now I can use (IC_M) to show the optimality of M's voting strategy. Suppose r_m is such that M votes yes in equilibrium, $\alpha(s_M, r_m) = 1$. Define $s^-_M := \max\{\widehat{s^-_M}|\widehat{s^-_M} < \psi(s_m) \leq s_M, s_m \in \mathcal{S}_m\}$. In other words, s^-_M is the signal just below the highest possible cutoff that can be induced by m so that M votes yes given s_M. The fact that M sometimes votes no ensures that s^-_M exists. Now, notice that due to monotonicity, switching from s_M to s^-_M only leads to a different allocation for messages in the set X. Any other message $r'_m \notin X$ must induce a cutoff either even smaller than s^-_M or larger than $\psi(r_m) = \psi(s_m)$, the cutoff induced by r_m. But then $\alpha(s_M, r'_m) = \alpha(s_M^-, r'_m)$. Therefore, I get:

$$\sum_{s_m \in \mathcal{S}_m} \Pr(s_m|s_M)\alpha(s_M, s_m)(\beta(s_m + s_M, 5) - c_M)$$
$$- \sum_{s_m \in \mathcal{S}_m} \Pr(s_m|s_M)\alpha(s_M^-, s_m)(\beta(s_m + s_M, 5))$$
$$= \sum_{s_m \in X} \Pr(s_m|s_M)(\beta(s_m + s_M, 5) - c_M) \geq 0$$

Now, (IC_M) implies that the first term of the difference is weakly larger than the second term. This inequality in turn implies that the difference is weakly positive. Finally, this result gives

$$\frac{\sum_{s_m \in X} \Pr(s_m|s_M)(\beta(s_m + s_M, 5) - c_M)}{\sum_{s_m \in X} \Pr(s_m|s_M)} \geq 0,$$

which states that the expected utility of M when voting yes conditional on r_m and s_M is greater than zero, the expected utility of voting no.

For any message r_m such that $\alpha(s_M, r_m) = 0$, the optimality of voting no can be shown analogously. Thus I have now established the optimality of M's voting strategy in equilibria in which she always votes yes, in equilibria in which she always votes no, and in all other equilibria. □

Lemma 2. Information transmission can only occur in equilibrium as long as $|q_M - q_m| \leq 1$.

Proof. Given Lemma 1, it is sufficient to consider equilibria in which m affects the outcome only through cheap talk and M affects the outcome only through voting. As stated in Corollary 3.1, m can use one of four different communication strategies in such an equilibrium.

Full information transmission is feasible only when $q_M = q_m$. Therefore what remains to be shown is that the two partial information transmission strategies cannot be sustained in equilibrium when $|q_M - q_m| > 1$. Then, the only strategy that remains for m is babbling, in which case there is no information transmission.

Suppose $q_M > q_m + 1$, for example $q_M = 4$ and $q_m = 2$. Consider strategy (a). After receiving $r_m = 0$, M only votes yes if $s_M \geq q_M$. After receiving $r_m \in \{1, 2\}$, M only votes yes ($a_M = 3$) either whenever $s_M + 1 \geq q_M$ or whenever $s_M + 2 \geq q_M$, depending on M's preferences. But then m would always want to deviate when

$s_m = 0$. In terms of the example, M never votes yes when $r_m = 0$ and either votes yes whenever $s_M \geq 3$ or whenever $s_M \geq 2$ when $r_m \in \{1, 2\}$. But then m never wants to send message $r_m = 0$. The same approach can be applied with strategy (b) or when $q_M + 1 < q_m$. Therefore, partial information transmission cannot occur in equilibrium. □

Lemma 3. In an equilibrium that implements a social choice α with partial information transmission and which is constructed with the strategies described in Lemma 1, the communication strategy of m depends on τ, q_M, and q_m in the following way.

With majority rule ($\tau = 3$):

- If $q_m < q_M$: m uses communication strategy (a).
- If $q_m > q_M$: m uses communication strategy (b).

With unanimity ($\tau = 5$):

- If $q_m < q_M$: m uses communication strategy (a).
- If $q_m > q_M$: Either communication strategy can occur.

Proof. As babbling and full information transmission are ruled out in equilibria with partial information transmission, Corollary 3.1 means that m must use either communication strategy (a) or (b). Therefore, to prove the lemma it suffices to show that in each of the first three cases that the other (i.e., the non-prescribed) communication strategy cannot be used in equilibrium.

If $q_m > q_M$, then $q_m = q_M + 1$ and if $q_m < q_M$, then $q_m = q_M - 1$.

- Majority rule, $q_m = q_M - 1$:

  Suppose m uses communication strategy (b). If $r_m = 2$, then M votes yes whenever $s_M \geq \psi_M(2) = q_M - 2$. Since there is partial information transmission, M votes differently if $r_m \neq 2$. But then, whenever $s_m = 1$, m would want to deviate and report $r_m = 2$, since M would vote yes whenever $s_M \geq \psi_M(2) = s_M - s_m - 1 = q_m - s_m$. In other words, M votes yes whenever $s_M + s_m \geq q_m$, a strategy which m prefers over any other possible voting strategy of M.

- Majority rule, $q_m = q_M + 1$:

  The argument that rules out communication strategy (a) is exactly analogous to the above argument.

- Unanimity, $q_m = q_M - 1$:

  Suppose m uses communication strategy (b). If $r_m = 2$, then M votes yes whenever $s_M \geq \psi_M(2) = q_M - 2$. Since there is partial information transmission, M votes differently if $r_m \neq 2$. But then, whenever $s_m = 1$, m would want to deviate and report $r_m = 2$, since M would vote yes whenever $s_M \geq \psi_M(2) = s_M - s_m - 1 = q_m - s_m$. In other words, M votes yes whenever $s_M + s_m \geq q_m$, a strategy which m prefers over any other possible voting strategy of M.

□

Proposition 3.2. If several social choice functions are implementable in equilibrium, there is always a Pareto-dominant one.

Proof. This proof is a case-by-case proof. Both for majority rule and for unanimity I show that whenever several social choice functions can be implemented in equilibrium, one social choice function Pareto-dominates the others.

Majority rule ($\tau = 3$)

One social choice function that is always implementable is the one that corresponds to no information transmission. With the strategies of Lemma 1 this social choice is implemented by m babbling and always voting $a_m = 2$ and M voting $a_M = 3$, whenever she prefers the reform solely based on her own signal s_M. Whenever information transmission is not feasible, this social choice function is the only implementable one, as no other social choice function without information transmission can be implemented with the strategies of Lemma 1.

If information transmission is feasible, then Lemma 2 means that $|q_M - q_m| \leq 1$. With $q_M = q_m$ the preferences of M and m are perfectly aligned. Therefore, the social choice function that corresponds to full information transmission clearly dominates any other social choice that is implementable.

Thus, what remains to be shown is that, when only partial information transmission is feasible, the social choice function corresponding to information transmission is preferred by both M and m to the social choice function corresponding to no information transmission. As m follows a different strategy depending on whether $q_M > q_m$ or $q_m < q_M$, I verify this statement for each case separately.

1. $q_M > q_m$, specifically $q_M = q_m + 1$.

   Lemma 3 gives that m follows communication strategy (a). Information transmission is not possible, whenever $q_m = 0$, because m would always want to lead M to vote pro reform, regardless of s_m. Hence $1 \leq q_m \leq 4$. Incentive compatibility gives that in the equilibrium with partial information transmission the reform is implemented if

$$s_M \geq q_M - \begin{cases} 0 & \text{if } r_m = 0 \\ 2 & \text{otherwise.} \end{cases}$$

   Since M can perfectly infer $s_m = 0$ if $r_m = 0$, she can only vote $a_M = 3$ if $s_M \geq q_M$. If $r_m > 0$, she would either want to vote for reform if $s_M \geq q_M - 1$ or $s_M \geq q_M - 2$. But in equilibrium it must be that she votes for reform only if $s_M \geq q_M - 2$, otherwise m would want to deviate to $r_m = 0$ whenever $s_m = 1$. With no information transmission the reform is implemented if $s_M \geq q_M - 1$, since $\beta(q_M, 5) = \beta(q_M - 1, 3) \geq c_M$.

   Call the partial information social choice function α and the no information transmission social choice $\widehat{\alpha}$. Now only consider $2 \geq q_M \leq 4$. These social choice functions differ only in cases: $\alpha(q_M - 2, 2) = 1$, $\alpha(q_M - 1, 0) = 0$, and $\alpha(q_M - 2, 1) = 1$, with $\widehat{\alpha}$ prescribing the opposite action.

   M prefers $\alpha(q_M - 2, 2) = 1$ and $\alpha(q_M - 1, 0) = 0$ to $\widehat{\alpha}$ but she prefers $\widehat{\alpha}(q_M - 2, 1) = 0$ to $\alpha(q_M - 2, 1) = 1$. Notice that M's ex-ante utility can be expressed as

$$\sum_{s_M \in \mathcal{S}_M} \Pr(s_M) \sum_{s_m \in \mathcal{S}_m} \Pr(s_m | s_M) \alpha(s_M, s_m)(\beta(s_m + s_M, 5) - c_M).$$

   Now comparing α to $\widehat{\alpha}$, M's ex-ante utility is higher with α since her utility gains from $\alpha(q_M - 2, 2) = 1$ outweigh her losses due to $\alpha(q_M - 2, 1) = 1$, since it is an equilibrium strategy for M to vote yes whenever $s_M = q_M - 2$

and $r_m > 0$, which implies

$$\Pr(s_m = 1|s_M = q_M - 2, r_m > 0)(\beta(q_M - 1, 5) - c_M) + \\ \Pr(s_m = 2|s_M = q_M - 2, r_m > 0)(\beta(q_M, 5) - c_M) \geq 0.$$

m prefers $\alpha(q_M - 2, 2) = 1$ and $\alpha(q_M - 2, 1) = 1$ to $\widehat{\alpha}$ but he prefers $\widehat{\alpha}(q_M - 1, 0) = 1$ to $\alpha(q_M - 1, 0) = 1$. However, similar to the argument for M, m's gains due to $\alpha(q_M - 2, 1) = 1$ outweigh his losses due to $\alpha(q_M - 1, 0) = 0$, since the ex-ante probability of $(s_M, s_m) = (q_M - 2, 1)$ is higher than the ex-ante probability of $(s_M, s_m) = (q_M - 1, 0)$ and $\beta(q_M - 2 + 1, 5) = \beta(q_M - 1 + 0, 5)$.

Finally, consider $q_M = 4$ and $q_m = 5$. Then the reform is never implemented in the equilibrium without information transmission. With information transmission, $\alpha(3, 1) = 1$ and $\alpha(3, 2) = 1$. Clearly, m prefers α to $\widehat{\alpha}$. M also prefers α since her gains from $\alpha(3, 2) = 1$ must outweigh her losses due to $\alpha(3, 1) = 1$ from an ex-ante perspective, as she votes yes with $r_m > 0$ and $s_M = 3$.

2. $q_M < q_m$, specifically $q_M = q_m - 1$.

The proof is similar to the proof for $q_M > q_m$. Lemma 3 gives that m follows communication strategy (b). Information transmission is not possible, whenever $q_M = 0$, because M always votes yes. Hence $1 \leq q_M \leq 4$. Incentive compatibility gives that in the equilibrium with partial information transmission the reform is implemented if

$$s_M \geq q_M - \begin{cases} 2 & \text{if } r_m = 2 \\ 0 & \text{otherwise.} \end{cases}$$

Since M can perfectly infer $s_m = 2$ if $r_m = 2$, she can only vote $a_M = 3$ if $s_M \geq q_M - 2$. If $r_m < 2$, she would either want to vote for reform if $s_M \geq q_M - 1$ or $s_M \geq q_M - 0$. But in equilibrium it must be that she votes for reform only if $s_M \geq q_M$, otherwise m would want to deviate to $r_m = 2$ whenever $s_m = 1$. With no information transmission the reform is implemented if $s_M \geq q_M - 1$, since $\beta(q_M, 5) = \beta(q_M - 1, 3) \geq c_M$.

Again, call the partial information social choice function α and the no information transmission social choice $\widehat{\alpha}$. These social choice functions differ only in cases: $\alpha(q_M - 2, 2) = 1$, $\alpha(q_M - 1, 0) = 0$, and $\alpha(q_M - 1, 1) = 0$, with $\widehat{\alpha}$ prescribing the opposite action.

M prefers $\alpha(q_M - 2, 2) = 1$ and $\alpha(q_M - 1, 0) = 0$ to $\widehat{\alpha}$ but she prefers $\widehat{\alpha}(q_M - 1, 1) = 1$ to $\alpha(q_M - 1, 1) = 0$. Again comparing α to $\widehat{\alpha}$, M's ex-ante utility is higher with α since her utility gains from $\alpha(q_M - 1, 0) = 0$ outweigh her losses due to $\alpha(q_M - 1, 1) = 0$, since it is an equilibrium strategy for M to vote no whenever $s_M = q_M - 1$ and $r_m < 2$, which implies

$$\Pr(s_m = 0 | s_M = q_M - 1, r_m < 2)(\beta(q_M - 1, 5) - c_M) + \\ \Pr(s_m = 1 | s_M = q_M - 1, r_m < 2)(\beta(q_M, 5) - c_M) \leq 0.$$

m prefers $\alpha(q_M - 1, 0) = 0$ and $\alpha(q_M - 1, 1) = 0$ to $\widehat{\alpha}$ but he prefers $\widehat{\alpha}(q_M - 2, 2) = 0$ to $\alpha(q_M - 2, 2) = 1$. However, similar to the argument for M, m's gains due to $\alpha(q_M - 1, 1) = 0$ outweigh his losses due to $\alpha(q_M - 2, 2) = 1$, since the ex-ante probability of $(s_M, s_m) = (q_M - 1, 1)$ is higher than the ex-ante probability of $(s_M, s_m) = (q_M - 2, 2)$.

Unanimity ($\tau = 5$)

The social choice function $\alpha(s_M, s_m) = 0$ for all s_M, s_m is always implementable with unanimity. In the corresponding equilibrium, no group has an incentive to deviate at the voting stage and m babbles. Due to (IR_M^1) and (IR_m^1), I know that any other implementable social choice function must Pareto-dominate this no reform social choice. With $q_M = q_m$, the social choice corresponding to full information transmission is implementable and again clearly Pareto-dominates any other social choice. For other social choice functions, the analysis again depends on whether $q_M > q_m$ or $q_M < q_m$.

1. $q_M > q_m$.

   There are three different cases in which α does not have information transmission, i.e., α is contingent on at most one group's signal. First, the social choice with the reform never being implemented. Second, when $q_M = 1$, then the reform is a social choice such that the reform is always implemented, since M always votes yes if m babbles, given $\beta(q_M, 5) = \beta(q_M - 1, 3)$. Third, if $q_M > 1$, then there is a social choice function corresponding to m babbling and always voting yes in equilibrium and M voting yes whenever $s_M > q_M - 1$. This social choice function, described for the majority rule case, is also feasible with unanimity. The reason is that m cannot meaningfully deviate at the communication stage and can only deviate at the voting stage by voting $a_m < 0$ for some s_m. But this deviation only matters if M votes yes, in which case $s_M \geq q_M - 1$, which implies that m would want the reform regardless of s_m. For $q_M > 1$, it is possible to have α only contingent on s_M and not s_m but not vice versa, since m can only always vote yes because $q_m < q_M$.

In any information transmission social choice function, Lemma 3 gives that m follows communication strategy (a). If both social choice functions are implementable, then this communication social choice dominates the no communication social choice. The underlying argument is exactly the same as the above argument with majority rule.

2. $q_M < q_m$.

Apart from the reform never being implemented and the reform always being implemented, the only social choice function without information transmission has α only depend on s_m, i.e., M always votes yes. This social choice function is feasible whenever c_m is such that $c_m > \beta(q_M - 1, 2)$. In this case, M never has an incentive to vote no, even if $s_M = 0$, since m already votes no whenever M does not want the reform. Again, this social choice function clearly Pareto-dominates never implementing the reform. However, when partial information transmission is feasible, it is not so clear that a Pareto-dominant social choice function must exist in the set of implementable social choice functions. With unanimity and $q_M + 1 = q_m$, unlike with unanimity and $q_M = q_m + 1$, Lemma 3 does not prescribe a communication strategy for m. Thus there are up to four different implementable social choice functions in this case. The social choice with which the reform is never implemented (1), a social choice function with $\alpha(s_M, s_m)$ only depending on s_m (2), an information transmission social choice function in which m uses communication strategy (a) in equilibrium (3) or one in which m uses communication strategy (b) in equilibrium (4). To show that one social choice function always Pareto-dominates other implementable social choice functions I investigate each possible case separately:

- $(q_M, q_m) = (0, 1)$

Apart from the reform never being implemented, the only other feasible social choice has the reform always implemented. The latter is feasible since M can clearly always vote yes and in this case m also wants to to so, since $\beta(0,2) > \beta(0,3) = \beta(1,5)$. Thus m can babble and always vote yes, even when $s_m = 0$.

- $(q_M, q_m) = (1,2)$

First, suppose that $\beta(0,2) \geq c_m$. Then, as with $(q_M, q_m) = (0,1)$ there exists a social choice function where both groups always vote yes. There can only be a social choice function in which m uses communication strategy (a), if the reform is never implemented whenever $s_m = 0$. Otherwise, when $r_m = 0$, M would vote yes whenever $s_M \geq 1$, but then m would want to deviate to $r_m = 0$ when $s_m = 1$. Now, let $\widetilde{\alpha}$ correspond to the reform always being implemented and $\widehat{\alpha}$ to the reform always being implemented unless $s_m = 0$. Notice that $\widetilde{\alpha}$ is not implementable given $\beta(0,2) \geq c_m$, as m would want to deviate whenever $s_m = 0$ given that

$$\beta(0,2) - c_m = \sum_{s_M \in \mathcal{S}_M} Pr(s_M | s_m = 0)(\beta(s_M, 5) - c_M) \geq 0.$$

Suppose there is also a social choice in which m uses communication strategy (b). Here it must be that the reform is always implemented when $r_m = 2$ and the reform is implemented when $s_M \geq 1$ when $r_m < 2$. Call this social choice α, which only differs from $\widetilde{\alpha}$ in that $\alpha(0,0) = \alpha(0,1) = 0$. Now clearly m prefers α to $\widetilde{\alpha}$. M prefers $\alpha(0,0) = 0$ and does not prefer $\alpha(0,1) = 0$. However, her gains due to $\alpha(0,0) = 0$ must outweigh her losses due to $\alpha(0,1) = 0$ since she votes no whenever $r_m < 2$ and $s_M = 0$.

Second, suppose that $\beta(1,2) \geq c_m > \beta(0,2)$. Now $\widetilde{\alpha}$, the social choice corresponding to always implementing the reform, does not exist. However, $\widehat{\alpha}$ is now implementable and clearly dominates never implementing the reform. If α, the information transmission social choice, is also implementable, it Pareto-dominates $\widehat{\alpha}$. α differs from $\widehat{\alpha}$ in that $\alpha(0,1) = 0$ and $\alpha(1,0) = \alpha(2,0) = \alpha(3,0) = 1$, while $\widehat{\alpha}$ has the opposite value in those cases. Now both M and m prefer $\alpha(2,0) = \alpha(3,0) = 1$. M prefers $\alpha(1,0) = 1$ to $\widehat{\alpha}$ but not $\alpha(0,1) = 0$ and vice versa for m. M's gains due to $\alpha(1,0) = 1$ outweigh her losses due to $\alpha(0,1) = 0$ since the ex-ante probability of $(s_M, s_m) = (1,0)$ is higher than the ex-ante probability of $(s_M, s_m) = (0,1)$. For m, his gains due to $\alpha(2,0) = \alpha(3,0) = 1$ outweigh his losses due to $\alpha(1,0) = 1$ since he votes yes when $s_m = 0$.

- $(q_M, q_m) = (2,3)$

  Now, since $c_M \geq \beta(1,5) = \beta(0,3)$, always implementing the reform is not a feasible social choice.

  First, suppose $\beta(1,2) > c_m$. Again, m can only utilize communication strategy (a) if the reform is never implemented whenever $s_m = 0$. When $s_m > 0$ and thus $r_m > 0$ the reform is either implemented always or implemented whenever $s_M \geq 1$, depending on c_M. Suppose the latter is the case and call this social choice $\widehat{\alpha}$, which dominates never implementing the reform. If there is also a feasible social choice in which m uses communication strategy (b) it must that when m reports $r_m = 2$ the reform is always implemented and when m reports $r_m < 2$ the reform is implemented when $s_M \geq 2$. Call this social choice α. It is not possible that when $s_m = 1$ and thus $r_m < 2$ the reform is implemented when $s_M \geq 1$ since m would want to deviate when $s_m = 2$. It is similarly not possible that the reform is never implemented when $s_m \leq 1$, since m would want to deviate to $r_m = 2$ as $\beta(1,2) > c_m$. Now, if both are implementable, α Pareto-dominates $\widehat{\alpha}$. The social choice functions differ in $\alpha(2,0) = \alpha(3,0) = 1$ and $\alpha(1,1) = 0$ with $\widehat{\alpha}$ having the opposite values. Both M and m prefer $\alpha(3,0) = 1$. M prefers $\alpha(2,0) = 1$ and does not prefer $\alpha(1,1) = 0$ to $\widehat{\alpha}$, while the opposite is true for m. For m, his gains due to $\alpha(3,0) = 1$ outweigh his losses due to $\alpha(2,0) = 1$, since he votes yes with $s_m = 0$. For M, her gains outweigh her losses since

  $$\begin{aligned}
  \Pr(s_M = 3, s_m = 0)(\beta(3,0) - c_M) \\
  + \Pr(s_M = 2, s_m = 0)(\beta(2,0) - c_M) \\
  &\geq 2\Pr(s_M = 2, s_m = 0)(\beta(2,0) - c_M) \\
  &= \Pr(s_M = 1, s_m = 1)(\beta(1,1) - c_M),
  \end{aligned}$$

  where the top row represents the utility gained with α compared to $\widehat{\alpha}$ and the bottom row represents the utility lost. Now suppose c_M is such that the reform is always implemented whenever m uses communication strategy (a) and $s_m > 0$. Call this social choice function $\widetilde{\alpha}$. Then the social choice functions differ in $\alpha(2,0) = \alpha(3,0) = 1$ and $\alpha(0,1) = \alpha(1,1) = 0$ with $\widetilde{\alpha}$ having the opposite values. Again clearly α Pareto-dominates $\widetilde{\alpha}$. The only difference to the previous case is that here $\alpha(0,1) = 0$ and $\widetilde{\alpha}(0,1) = 1$. But both M and m prefer $\alpha(0,1) = 0$.

Second, suppose $\beta(1,2) < c_m$. Now, $\widetilde{\alpha}$ is not implementable, as m would want to deviate to $r_m = 0$ when $s_m = 1$. However, now a social choice function is implementable in which m uses communication strategy (b), the reform is never implemented when $s_m \leq 1$ and always implemented when $s_m = 2$. Call this social choice $\overline{\alpha}$. If α is implementable as well, α Pareto-dominates $\overline{\alpha}$. The two allocations differ in that $\alpha(1,1) = \alpha(2,1) = \alpha(3,1) = 1$. Clearly M prefers α. m must also prefer α, given that he must vote yes whenever $s_m = 1$. Thus his gains from $\alpha(2,1) = \alpha(3,1) = 1$ must outweigh his losses due to $\alpha(1,1) = 0$.

- $(q_M, q_m) = (3, 4)$

Suppose m follows communication strategy (a). Then it must be that the reform is never implemented when $s_m = 0$. Similarly, it must be that the reform is only implemented when $s_m \geq 1$ and $s_M \geq 2$. There can be no equilibrium in which M votes yes whenever $s_M \geq 1$ when receiving $r_m \neq 0$. The reason is that m's posterior of the success of the reform conditional on $s_M \geq 1$ and $s_m = 1$ is smaller than c_m as

$$\Pr(s = 1 | s_m = 1, s_M > 0) = \frac{1 - (1-p)^3}{1 - (1-p)^3 + 1 - p^3} < p = \beta(3,5) < c_m$$

for all $p \in (0.5, 1)$. Therefore, if M voted yes whenever $r_m \neq 0$ and $s_M \geq 1$, m would have an incentive to deviate and vote no when $s_m = 1$. Call this social choice α.

Suppose m follows communication strategy (b). Then it must be that reform is never implemented when $s_m \leq 1$, since m never wants the reform if $s_m = 0$. If $s_m = r_m = 2$ the reform is implemented whenever $s_M \geq 1$. Call this social choice $\widetilde{\alpha}$.

If only one of the two social choice functions α and $\widehat{\alpha}$ is implementable, it clearly Pareto-dominates the only other implementable social choice function, never implementing the reform. If both α and $\widehat{\alpha}$ are implementable, α dominates. The two functions only differ in that $\alpha(2,1) = \alpha(3,1) = 1$ and $\alpha(1,2) = 0$. m prefers $\alpha(1,2) = 0$ and $\alpha(3,1) = 1$ but not $\alpha(2,1) = 1$. However, his gains due to $\alpha(3,1) = 1$ outweigh his losses due to $\alpha(2,1) = 1$ as he votes yes whenever $s_m = 1$. M prefers $\alpha(2,1) = 1$ and $\alpha(3,1) = 1$ but not $\alpha(1,2) = 0$. However, her gains due to $\alpha(2,1) = 1$ outweigh her losses due to $\alpha(1,2) = 0$ since $(s_M, s_m) = (2,1)$ is twice as likely as $(s_M, s_m) = (1,2)$ from an ex-ante perspective.

- $(q_M, q_m) = (4,5)$

  In any social choice for which the reform is sometimes implemented, m can only follow communication strategy (b). If such a social choice is implementable, the reform is only implemented when $s_m = 2$ and $s_M \geq 2$.

- $(q_M, q_m) >> (4,5)$

  There is only one implementable social choice function. The reform is never implemented.

$\square$

4. In or Out - the effect of small parties winning representation in proportional representation systems. A regression discontinuity design.

Introduction

In legislative bodies with a proportional representation electoral system there are often small parties that fail to win representation. This failure can be caused by a minimum electoral threshold, or by the overall number of seats up for election being too small for a small vote share to translate to at least one seat. In this paper, I estimate the causal effect of parties winning representation at the extensive margin, i.e., the effect of being in as opposed to being out. Identification of a causal effect is a tricky issue, as winning representation is not random. Parties that win representation are inherently different from parties that do not, and not only in the sense that the former gain more votes.

To address this issue of causal identification, a lot of attention in empirical Economics has been given to quasi-experimental research designs. A prominent example of such a research design is the regression discontinuity (RD) design. It exploits settings in which treatment is assigned based on a score relative to a cut-off. The idea underlying this approach is to consider treatment assignment for observations with a score close to the threshold as essentially random. The RD approach has been particularly well received in Political Economy. Following the seminal work of Lee (2008), who estimates the incumbency advantage[1] in close US House elections, many researchers have used an RD approach to exploit close elections as natural experiments. Almost all of these studies are based on the threshold of winning at 50 percent in a majoritarian electoral system with two parties or two candidates.

In proportional representation electoral systems there is no fixed threshold of winning at 50 percent. Nevertheless, there are discontinuities stemming from the seat allocation method that can be exploited for an RD design. Folke (2014) is the first and only one, to my knowledge and not counting this paper, to build an RD framework based on these discontinuities. He studies the effect of certain parties winning representation on the intensive margin on policy, i.e., the effect of winning an additional seat. In this paper, in contrast to his approach, I develop a similar RD setup to study the effect of any party winning representation at the extensive margin, i.e., of winning at least one seat. Winning representation at the extensive margin potentially has a two-dimensional effect: First, an effect on policy, through the effect on the overall composition of the municipal council. Second, an effect on the party itself, such as an incumbency effect similar to the one estimated by Lee (2008) in a majoritarian context.

With respect to the first dimension, I can investigate the relationship between fractionalization and public spending. A party winning representation, in most cases, increases the total number of parties in the legislature by 1 and thus increases the fractionalization in the legislature. The Political Economy literature generally predicts a positive relationship between fractionalization and spending, either within an electoral system (Lizzeri and Persico, 2005), or with respect to the dichotomy of a majoritarian system with two parties and a proportional representation system with many parties (Milesi-Ferretti, Perotti, and Rostagno, 2002; Persson, Roland, and Tabellini, 2007).

With respect to the second dimension I can add to the Political Economy literature on incumbency effects. The existence of incumbency effects is a well-established fact for majoritarian electoral systems in many different contexts.[2] The effect of

[1]The average advantage an incumbent has over a challenger in terms of the vote share.

[2] Apart from the studies with an RD approach, such as Lee (2008), there are many more studies with different methodological approaches that investigate incumbency effects. For example, Gelman and King (1990), Alford and Brady (1993), and Levitt and Wolfram (1997) are are some influential studies on incumbency effects in the US House not using an RD approach.

incumbency in a proportional representation system, i.e., the effect of winning representation, has been studied much less by comparison.

For my estimation, I use data from municipal council elections in the German state of Thuringia. Using the seat allocation method in these elections, I calculate the counterfactual number of votes a party would have needed to safely win representation in the municipal council. I use the difference in the actual number of votes to this threshold as a measure of closeness of the election for a party. I estimate a sharp RD design, where the discontinuous jump in an outcome can be interpreted as the average causal effect of winning representation.

For public spending, measured by investment spending, I find a large and statistically significant effect of a party winning representation and thus increasing the overall number of parties in the municipal council by 1. With the additional party, investment spending per capita per year increases by 142 Euros, representing an increase of almost 50 percent. For the effect of incumbency on the individual party, I find that small parties close to the extensive margin have a high probability of dropping out and not running in the next election. However, I find a statistically significant difference suggesting that parties that win seats are roughly 35 percent less likely to drop out in the next election. This phenomenon precludes the exact estimation of an incumbency effect with respect to the vote share. I only observe the vote share in the next election for a selected sample and selection varies depending on whether a party won seats or not. In a bounding exercise, I estimate a lower bound and an upper bound for the incumbency effect. Both bounds are relatively close to zero and not statistically significant. Therefore, I do not find evidence of a vote share incumbency effect.

With respect to the estimated significant discontinuities for investment spending and dropping out, I also attempt to disentangle the estimated average effect and thus to identify potential channels of the effect. One potential channel is the effect of winning representation on the composition of the municipal council. If a party winning representation causes a larger change in the composition of the municipal council, I expect a larger effect on outcomes as well. To test this prediction, I interact the jump in treatment with a measure of the impact of the party close to the threshold on the composition of the municipal council. As measures I use both the change in the Herfindahl index attributable to that party and the baseline number of parties that win seats, not counting that party. If the baseline number of parties is small, one additional party represents a larger change. For investment spending, interaction effects are significant and have the expected sign, suggesting that a greater impact on the composition of the municipal council is associated with a larger effect on spending. For the probability of dropping out, I do not find statistically significant interaction effects.

Overall, this paper contributes to the Political Economy literature by shedding light on the effects of incumbency at the extensive margin. Compared to majoritarian electoral systems, these effects are much less studied for proportional representation systems. However, knowledge about them can help policy makers to make informed decisions. For example, with respect to constitutional design, policy makers can better answer whether an electoral system should include electoral thresholds. Electoral thresholds are put in place to prevent very small parties from gaining seats and thus to prevent the number of parties that win representation from becoming too large. My results suggest a considerable effect of an additional party winning representation on political outcomes.

Related literature

The RD approach has been applied frequently in Political Economy with respect to many different questions, mostly exploiting the discontinuity in close elections. Tables listing some of the these works can be found both in Lee and Lemieux (2010) and Caughey and Sekhon (2011).

The vast majority of works that apply the RD approach to close elections exploit the 50-50 threshold. In majoritarian two-party systems, with only one seat in contention within a district, this threshold of absolute majority is in fact the only relevant discontinuity. In proportional representation systems with many parties, there is no clear 50-50 threshold. In the context of municipal council elections in Sweden, Pettersson-Lidbom (2008) and Liang (2013) overcome this obstacle by lumping political parties into two-camps: left-wing and right-wing. This way, they obtain a 50-50 threshold for overall control of the municipal council, despite municipal council elections in Sweden having a proportional representation system and generally more than two parties.

Instead of having a natural discontinuity at the 50-50 threshold, proportional representation systems induce multiple discontinuities that can be exploited. This feature is due to the fact that a continuous vote share is translated into a discrete number of seats. Also in the context of Swedish municipal council elections, Carlsson, Dahl, and Rooth (2015) exploit this fact for a control-function approach, which relies on the same discontinuities for identification that RD studies use but is methodically different from the RD approach. Besides applying the RD approach to the 50-50 threshold, Liang (2013) also uses a control-function approach. To my knowledge, Folke (2014) was the first to develop an RD methodology specifically for the proportional representation context. Again, using Swedish municipal council elections, he investigates how an additional seat for a certain party in a municipality impacts policy. His approach to identifying the relevant thresholds is methodically different, due to the different seat allocation rules in Sweden and Thuringia, but conceptually very similar. One crucial difference of Carlsson, Dahl, and Rooth (2015), Liang (2013), and Folke (2014) to this paper is that they predominately consider the intensive margin while I consider the extensive margin. They focus on the effect of additional seats for a party while I look at the effect of that party winning either some seats or no seats at all.

In the study that started the surge in the application of the RD approach to elections, Lee (2008) finds a strong incumbency effect. Being the incumbent comes with an advantage over the challenger in terms of votes in the next election.[3] He also finds that incumbent candidates are much more likely to run in the next election than candidates that narrowly lost the last one. Trounstine (2011) finds similar effects in the context of mayoral elections in California. However, they encounter this phenomenon in the two-party majoritarian context of the US. If the Democratic candidate decides not to run again, another candidate will take his or her place. The effect I find in the proportional representation setup is with respect to the party dropping out in the next election. Thus the aforementioned studies do not encounter the same sample selection problems I face in the estimation of an incumbency effect with respect to the vote share in the next election.

For the proportional representation context, Liang (2013) finds no evidence of a "ruling effect", an effect of belonging to the majority camp, on the vote share in the next election, but he finds an intensive margin incumbency effect. Additional seats on the municipal council lead to a higher vote share in the next election. Carlsson, Dahl, and Rooth (2015) also find such an effect and link it to the impact of strength in the municipal council on attitudes. Apart from these intensive margin effects,

[3]This advantage has been observed independently in many different contexts. Just for the US, there are studies which, using an RD design, find an incumbency advantage in House elections (Lee, 2008), state legislature elections (Uppal, 2010), and mayoral elections (Ferreira and Gyourko, 2009; Trounstine, 2011).

using only parties that win 0 seats or 1 seat, Liang (2013) also finds an extensive margin effect, i.e., the effect of winning representation at all. However, in both papers the authors do not report a dropping out effect, as observed by Lee (2008) and Trounstine (2011) in a majoritarian context and in this paper in a proportional representation context. Consequently, Liang (2013) does not face a selection bias issue in estimating the extensive margin incumbency effect.

There are several studies, both theoretical and empirical, according to which having a proportional representation rather than a majoritarian electoral system is associated with more public spending. Persson and Tabellini (2004) provide an overview of these studies. Milesi-Ferretti, Perotti, and Rostagno (2002) develop a model of electoral competition that predicts more targeted and inefficient rather than diffused spending in proportional representation systems compared to majoritarian systems. They also provide empirical evidence in support of their model using a cross-country panel of OECD and Latin American countries. Persson, Roland, and Tabellini (2007), on the other hand, propose a model that attributes the difference between the systems to the higher incidence of coalition governments under proportional representation. Also using a cross-country panel, they give some empirical evidence in support of their assertion. Finally, instead of comparing the two electoral modes, Lizzeri and Persico (2005) provide a model of electoral competition within a proportional representation system that predicts more targeted spending with a larger number of parties.

Some researchers have used the RD approach to investigate public spending. Ferreira and Gyourko (2009) look at mayoral elections in the US and find no effect of the mayor's party on the amount of public spending. Albouy (2013) finds that members of Congress in the US attract more federal spending to their district if they belong to the party currently holding the majority. Republicans and Democrats also differ in the type of spending they attract. The former attract more spending for defense and transportation, the latter attract more spending for education and urban development. Pettersson-Lidbom (2008) finds a different effect of left-wing and right-wing governments on spending. Left-wing governments are found to spend more than right-wing governments. Again with regard to municipal councils in Sweden, Pettersson-Lidbom (2012) does not look at elections but exploits the thresholds that determine the size of the municipal council by population size. He finds that larger municipal councils lead to smaller government and explains this finding with the control function performed by the municipal council.

In any application of the RD approach, it is prudent to investigate whether the identifying assumption holds, i.e., if treatment is essentially random for observations close to the threshold. A large difference in pre-treatment observables or in the mass of observations on either side of the threshold might indicate discontinuous sorting, which would imply that the identifying assumption is violated. Caughey and Sekhon (2011) raised this concern in the context of congressional elections in the United States. They point out that even when investigating a narrow window around the 50-50 threshold close races might not be random. For example, they find that incumbents are more likely to win close elections than challengers. Snyder, Folke, and Hirano (2015) address this finding. They point out that this phenomenon is not necessarily indicative of manipulation. Instead, it can easily arise as a result of electoral competition and the empirical distribution of voting results. Eggers, Fowler, Hainmueller, Hall, and Snyder (2015) test the validity of the RD setup in many different election contexts. In general they fail to find problems due to sorting and argue that the concern raised by Caughey and Sekhon (2011) may be a fluke. Both Caughey and Sekhon (2011) and Eggers, Fowler, Hainmueller, Hall, and Snyder (2015) include recommendations for tests of the identifying assumption which guide my robustness tests. Among those tests, the test for a discontinuity in the histogram of the forcing variable was first proposed by McCrary (2008).

In applying the RD methodology, I am guided largely by Imbens and Lemieux (2008) and Lee and Lemieux (2010). For bandwidth choice in local linear regression I rely on Imbens and Kalyanaraman (2012).

The RD approach

In this section, I briefly describe the regression discontinuity (RD) approach in general and how I apply it in this context. An RD design is a quasi-experimental research method to identify the causal effect of a treatment. RD approaches can be applied in settings in which treatment status is not assigned randomly - as it would be in a controlled experiment - but assigned according to some characteristic, called the forcing variable, in relation to a threshold.

In this case, I am interested in the effect of party i making it into the municipal council. The treatment D_i is defined as

$$\begin{aligned} D_i = 1 \qquad & \text{if party } i \text{ holds at least one seat,} \\ D_i = 0 \qquad & \text{if party } i \text{ holds no seat,} \end{aligned}$$

and the treatment effect on outcome Y_i is $Y_i(D_i = 1) - Y_i(D_i = 0)$. Unfortunately, it is generally impossible to observe $Y_i(D_i = 1)$ and $Y_i(D_i = 0)$ for i at the same time. Also, unlike in an experiment, one cannot simply divide the sample into subgroups by treatment and compare outcomes to identify the treatment effect. Parties with $D_i = 1$ gain a larger number of votes than parties with $D_i = 0$ due in part to systematical differences between both groups. Consequently, differences in outcomes cannot only be attributed to the effect of the difference in the treatment D_i.

An elegant way to address this problem is the RD approach. The underlying idea is to exploit randomness in the realization of the forcing variable, so that treatment status can be considered random for observations that lie close to the threshold. Applied to this setting, whether $D_i = 1$ or $D_i = 0$ is close to random if the number of votes for party i is in the neighborhood of the number of votes i would need to win at least one seat. Consequently, parties that narrowly win seats are comparable to parties that narrowly do not and differences between these two groups can be attributed to the difference in D_i. By comparing two groups, it is possible to obtain a measure of the average treatment effect for these parties, $E[Y_i(D_i = 1) - Y_i(D_i = 0)|i \text{ close to threshold}]$.

The crucial methodological step in applying the RD approach in a proportional representation setting is to construct a one-dimensional forcing variable that quantifies the closeness of an election outcome. Folke (2014) was the first to develop an approach for the Sainte-Laguë seat allocation method. I use a conceptually slightly different approach for the seat allocation method described in Section III.i, which is based on the largest remainder method.

First, I calculate the counterfactual variable $\widehat{v_i}$, the number of votes party i would have needed to safely (without coin toss) win at least one seat, given that the number of votes for all other parties had stayed the same. The difference between party i's actual number of votes v_i and $\widehat{v_i}$ then measures how narrowly i won seats or missed winning seats in terms of the number of votes. To construct a measure comparable across different municipalities of different sizes, I divide the difference by the sum of votes for all other parties plus $\widehat{v_i}$, the counterfactual number of votes of party i. The resulting forcing variable x_i is a normalization of the vote share of party i centered at the threshold and can therefore be interpreted as measuring closeness in percent,

$$x_i = \frac{v_i - \widehat{v_i}}{\sum_{j \neq i} v_j + \widehat{v_i}} \cdot 100. \tag{4.1}$$

Given the construction of x_i, the relevant threshold determining treatment status is at 0,

$$D_i = \begin{cases} 1 & \text{if } x_i \geq 0 \\ 0 & \text{if } x_i < 0. \end{cases} \tag{4.2}$$

The crucial difference between my forcing variable and the forcing variable proposed by Folke (2014) is that I only look at a vote change for party i, keeping the number of votes for all other parties equal, while he considers a vote change in the entire vector of votes. For a detailed explanation why my much simpler approach is appropriate for the largest remainder method but not for the Sainte-Laguë method, see the Online Appendix of Folke (2014).

Specification

To estimate an RD design, I need to regress the outcome of interest on the forcing variable allowing for a discontinuous jump at the threshold. This jump represents the treatment effect. When using a parametric specification, RD estimates can be very sensitive to a misspecification of the functional form. One possible remedy of this problem is to use a semi-parametric estimation technique, such as local linear regression. I follow this approach for the main analyses and report alternative polynomial specifications as a robustness check.

In local linear regression at a point c, observations are weighted by closeness to the point with a Kernel function. As recommended by Imbens and Kalyanaraman (2012), I use the triangular Kernel,

$$K(x_i, c, h) = 1 - \left| \frac{x_i - c}{h} \right|. \tag{4.3}$$

In the RD approach with local linear regression, how many observations around the threshold are included in the estimation and how they are weighted is governed by the choice of bandwidth h. The trade-off in selecting the bandwidth is between identification and variance. Including observations further away from the threshold reduces the variance of the estimator but makes the identification assumption, that treatment status can be considered random, less plausible. I calculate the bandwidth according to a method proposed by Imbens and Kalyanaraman (2012), which yields a bandwidth that is asymptotically optimal under squared error loss.

Essentially, local linear regression can be described as a local and weighted version of a linear least squares regression. With Y_i as the outcome of interest, the regression is estimated optimizing

$$\min_{\alpha,\tau,\beta,\gamma} \sum_{\{i:-h\leq x_i\leq h\}} (Y_i-\alpha-\tau D_i-\beta x_i-\gamma x_i D_i)^2 K(x_i,0,h). \tag{4.4}$$

The effect of interest, the treatment effect, is captured by the discontinuity τ which represents the shift of the intercept at the threshold. The intercept α, the slope parameter β, and the change of the slope γ for treated observations control for the relationship between forcing variable and outcome and are not interpreted.

Given the identification strategy of the RD approach it is not necessary to include control variables in the specification. While an outcome of interest may be influenced by many different variables, this effect biases $\widehat{\tau}$, the estimate of τ, only if it itself is discontinuous at the threshold. In this case, however, the identifying assumption, that observations slightly below and slightly above the threshold are similar, is likely to be violated. Nevertheless, it is possible to include controls by estimating

$$\min_{\alpha,\tau,\beta,\gamma} \sum_{\{i:-h\leq x_i\leq h\}} (Y_i-\alpha-\tau D_i-\beta x_i-\gamma x_i D_i-\delta' Z_i)^2 K(x_i,0,h), \tag{4.5}$$

where Z_i represents a vector of controls. Imbens and Kalyanaraman (2012) also provide a method for bandwidth calculation adjusted for control variables, which I use.

There are two reasons why it may be useful to include controls. First, including controls can reduce the variance of the estimator. Second, including controls can also work as a basic check of the identifying assumption. If the inclusion of controls has a large impact on the result of estimation, it casts doubt on the identifying assumption.

Municipal politics and data

The lowest level of geographical administrative division in Germany is the level of the municipality (*Gemeinde*). I use municipal council elections in the East German state of Thuringia in 1994, 1999, and 2004.[4] In 1994, municipal elections were held in 1247 municipalities in Thuringia. These municipalities are very heterogeneous in terms of size, ranging from 47 inhabitants to 213,472 inhabitants.[5]

Municipalities in Germany perform a wide variety of functions. Among other functions these typically include the maintaining of schools and kindergartens, waste disposal, the provision of local infrastructure, emergency management, and the provision of local facilities for cultural and sporting events. To finance these functions, municipalities receive funding from higher level administrative units, e.g., the state, charge fees for some of their services, e.g., waste disposal, and levy certain taxes. Most importantly, in Germany the property tax (*Grundsteuer*) and the business tax (*Gewerbesteuer*) are levied at the municipal level.

The voting method

Municipal councils in Thuringia are elected using a party list proportional representation system.[6] Every voter gets three votes which he or she can freely allocate among the party lists running in the election. A seat allocation method is then used to translate continuous vote shares into discrete numbers of seats. The method applied in the elections in the data is the largest remainder method, also called Hare-Niemeyer, or Hamilton method. In this method, in a first step, seats are allocated according to the integer part of the vote share times the total number of available seats. In the second step, the remaining seats are allocated in the order of the size of the remainders.

Additionally, two adjustments are made to this basic largest remainder method in the elections in the data. A majority clause is added to ensure that any party

[4]Election data for the elections in 2009 and 2014 is also available. However, from 2009 on the election mode did not include a 5 percent minimum threshold. As this change affects the threshold of necessary votes to win seats and the composition of parties at the threshold, I do not include the 2009 election round. For 2014, outcome data is not yet available.

[5]Larger cities such as the capital Erfurt with 213,472 inhabitants are classified as *Kreisfreie Städte*. As such, they perform the functions of a municipality and the next highest unit of administration, the district (*Kreis*). I include them in my dataset.

[6]In some municipalities a majoritarian voting system was used instead. I address this issue in Section III.iv.

Figure 4.1: A flowchart describing the largest remainder seat allocation method with 5% threshold and majority clause.

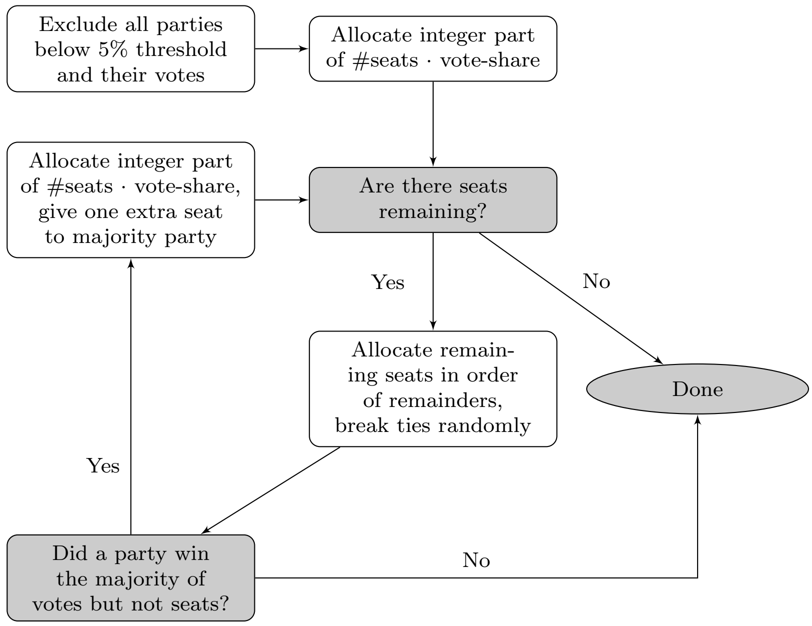

that wins a majority of votes also wins a majority of seats. Also, a 5 percent minimum threshold prevents any party with a vote share below 5 percent from winning any seats. For larger municipalities, this threshold implies that a party crossing it might go from no seats to several seats instead of going from no seats to one seat. Consequently, I do not refer to the effect at the extensive margin as the effect of winning the first seat. The resulting seat allocation method is depicted in a flowchart in Figure 4.1.

Politics on the municipal level

Politics on the federal or on the state level in Germany are dominated by major political parties, such as the center-right Christian Democrats (CDU), the center-left Social Democrats (SPD), the business-friendly Liberal Party (FDP), or the Greens. On the municipal level this statement does not hold true. While the traditional political parties also play a role, in many municipalities in Thuringia politics are dominated by free voter groups and local clubs.

Free voter groups are associations of local voters that enter municipal elections as a party list without formally having the legal status of a political party. They often represent specific local interests that are unrelated to the ideological differences between the major political parties. Chief among those voter groups are those formed by local clubs, such as local sports clubs, volunteer firefighters, community gardening clubs, or clubs for the preservation of local heritage. In addition, there are voter groups for the local church, or the PTA of the local school or local kindergarten.

Clubs such as the local football club or the volunteer firefighter club are an important part of the community, especially in more rural areas. Also, they are often directly affected by municipal politics. For example, the municipality of Marksuhl planned to spend 18,000 Euros on refurbishing local football pitchtes, 3,000 Euros on equipment for local volunteer firefighters, and 2,000 Euros on equipment for the kindergarten.[7] Therefore, local clubs often have a direct interest in participating in municipal politics.

The composition of parties running in a municipality is not necessarily stable over time. Some parties drop out, some parties merge and run as a joint party list, or parties that have previously formed a joint list split up again. The phenomenon can be observed for free voter groups and major political parties alike. In some cases major parties merge with free voter groups, e.g., one observation in the data is the party list “Social Democrats and Volunteer Firefighters” winning one seat on the municipal council of Oppurg in 2004.

Another major difference between municipal and state or federal level politics in Germany is that the municipal council does not appoint the government of the municipality. The mayor of a municipality is directly elected and does not need the continuous support of a majority of the municipal council. Therefore, parties in a municipal council do not formally form coalitions but cooperate on an issue-by-issue basis to reach a majority. Nevertheless, parties with similar political preferences often cooperate with each other, thus forming a type of informal coalition.

Main outcomes

I am interested in the effect of the treatment on municipal politics, i.e., what is the effect of a party narrowly winning or not winning representation in the municipal council. In particular, I want to answer two questions. First, what is

[7] According to the municipal gazette of Marksuhl (“Marksuhler Nachrichten”, 1/2006, “Gemeinsames Amtsblatt der Gemeinden Marksuhl, Wolfsburg-Unkeroda und Ettenhausen an der Suhl”).

the treatment effect through the resulting change in the number of parties in the municipal council on policy in the municipality? Second, what is the effect of holding or not holding seats on the party in question?

For the first question, I investigate realized investment spending between election rounds. There are several reasons for choosing spending, and in particular investment spending, as a policy measure. There is much more variation in the spending level than in other policy dimensions, such as, for example, tax rates, which are quite stable over time. Furthermore, as previously laid out, there is an existing theoretical Political Economy literature on the relationship between the number of represented parties and spending from which I can derive predictions. I have picked investment spending rather than overall spending, because it is more likely to reflect the policy of the current municipal council. For example, overall spending contains wage payments, debt repayment, and accounting devices such as internal loan repayment, which are likely to be directly caused by past policies rather than policies of the current municipal council. In addition, the type of investment spending mentioned in Section III.ii is a type of targetable policy that is likely to be affected by the size and composition of the municipal council. An additional party representing specific local interests that makes it onto the municipal council could lead to additional investment spending in line with these interests.

For the second question, I look at a dummy variable of whether the party in question drops out and does not run in the next election ($t + 1$). Parties that merged with another party are not counted as having dropped out. This measure is similar to the measures of Lee (2008) and Trounstine (2011) in the two-party majoritarian context, who instead look at candidates who do not run again for their party in the next election. Additionally, it is an indicator of the effect of winning representation on the success of the party in municipal politics.

For parties that do run in the next election, I can ask whether and how winning representation affects the success of the party in terms of votes. Unfortunately, the effect of the treatment on the propensity to drop out in the next election precludes the exact estimation of such an incumbency effect. The fact that parties with seats on the municipal council have a different propensity to drop out in the next election than parties without seats is likely to lead to a selected sample - as the decision to drop out is probably also related to the potential to attract votes. A party that decides to run again despite having narrowly missed winning seats in the last election is on average likely to be more optimistic about its prospects compared to a party that decides to run and that has narrowly won seats.

I estimate an upper bound on the incumbency effect by assigning the value 0 for the vote share in the next election to parties that drop out. Despite having dropped out, those parties most likely would have received some votes. Given the higher rate of dropping out to the left to the threshold, this procedure produces an upward biased estimate of the discontinuity. Similarly, I estimate a lower bound on the incumbency effect by substituting the missing vote share in the next election with the current vote share for parties that drop out. Parties that drop out despite having won representation might do so because they expect to do worse in the next election. Parties that drop out after not winning representation, on the other hand, might do so just because they did not win. As a result, the estimate of the discontinuity would be biased downwards.

Constructing the dataset

All data I use is publically available at the statistical office of the state of Thuringia (Thüringer Landesamt fÃ¼r Statistik)[8]. I use municipal election data for the 1994, 1999, and 2004 election rounds, census data from 1994 to 2008, and data on municipality finances from 1995 to 2008. For municipality spending outcomes, I use the average per capita spending in the years between election periods. For example, for an observation from the 1994 election round I investigate the average of per capita investment spending in 1995, 1996, 1997, and 1998.

In some municipalities, only one party list or only individual candidates registered for the election. In these cases, a majoritarian voting system was used instead of proportional representation. I exclude these municipalities.

Moreover, throughout the nineties, municipalities in Thuringia were frequently reorganized. Most of them merged to form larger municipalities. I exclude all observations in which the relevant municipality was reorganized prior to the next municipal election round.

Given the idea underlying the identifying assumption, the level of observation is the party. In other words, one observation corresponds to one of potentially many

[8] www.statistik.thueringen.de

party lists close to the threshold in a given municipality. However, one of my two main outcomes, investment spending, is aggregated at the level of the municipality.

To allow me to estimate municipality effects, I construct the dataset such that one observation corresponds to one municipality in an election round. To this end I only pick the party closest to the threshold in any municipality. In most cases, this party increments the total number of parties with seats on the municipal council by 1 by narrowly winning or not winning seats. In other words, if the party closest to the threshold narrowly won seats, it took these marginal seats from larger parties that still remained on the municipal council. In a few cases, 22 in total, the party closest to the threshold simply replaced or would have replaced another party compared to the counterfactual and thus did not affect the overall number of parties with seats. By excluding these cases, I am able to estimate a sharp RD design in which I can interpret treatment status as having an additional party on the municipal council. Put differently,

$$\begin{aligned} D_i = 1 \qquad & b+1 \text{ parties with seats,} \\ D_i = 0 \qquad & b \text{ parties with seats,} \end{aligned}$$

where b represents the baseline number of parties winning seats, without counting the party closest to the threshold.

For the final dataset, I then take all observations with a value of the forcing variable that is "sufficiently close" to the threshold. Naturally, there is a lower bound on the distance to the threshold. If 5 percent of the vote would have been enough to win representation, then $x_i \approx -5$. As most parties with $x_i < 0$ are much closer to winning seats, I define "sufficiently close" as values x_i close enough to 0 so that the support of x_i exhibits no gaps. To this end, I pick the top 90 percent of observations with $x_i < 0$. For positive values of x_i, I then use the same cutoff in absolute terms.

Finally, there are some cases in which I cannot reproduce the reported seat allocation from the election result. This problem might occur due to data entry mistakes, the seat allocation might have been calculated wrongly, parties that won seats might have refused to fill them, or the election result might have required a coin toss. As this problem affects only a small number of observations (7 in total), I exclude them.

Summary statistics of the final dataset can be found in Table 4.8.

Trimming with respect to investment per capita

One of my two main outcomes of interest is investment spending per capita in the municipality. The distribution of investment contains some large outliers. While the average is 349 Euros, the municipality of Petersberg spent on average 7092 Euro per capita in 1995-1998. These large outliers pose a threat to my estimation. To address this problem, I trim the data, trimming off the upper tail of the distribution. For my main set of results, I trim the data at the 95th percentile. I discuss the sensitivity of my results to this procedure in Section V.i.

Results

At the beginning of this section, I give my main set of empirical results. The following subsections contain an investigation into the possible channels of the treatment effect, in relation to the existing literature.

Table 4.1 reports the results for the number of parties that win seats in the municipal council. My results for my main outcomes of interest, investment per capita, dropping out in the next election and the bounds on the incumbency effect, are presented in Tables 4.2, 4.3, 4.4, and 4.5. Graphically, the discontinuity for these outcome variables are depicted in Figures 4.2 and 4.3 for the baseline estimation. In all three tables, the first column contains the baseline. In columns 2-4 either population size, or an election round dummy, or both are added as control variables. In columns 5 and 6 the baseline specification is estimated with half and twice the optimal bandwidth for comparison.

The results of Table 4.1 provide a first sanity check of the estimation strategy. As I am interested in the effect of the number of parties in the municipal council on investment spending, the dataset is constructed so that in the counterfactual case where party i landed just on the other side of the threshold, the number of parties in the municipal council would change by exactly 1 party, party i. Put differently, for Y_i being the number of parties winning seats, it is possible to know both $Y_i(D_i = 1)$ and $Y_i(D_i = 0)$ at the same time due to the seat allocation rule. For the treatment effect we have $Y_i(D_i = 1) - Y_i(D_i = 0) = 1$ by construction.

FIGURE 4.2: Local linear regression conducted separately on both sides of the threshold. Outcome: The number of parties winning seats in the election and (trimmed) investment spending. The box in the bottom right corner of each panel contains the estimate of the discontinuity (treatment effect).

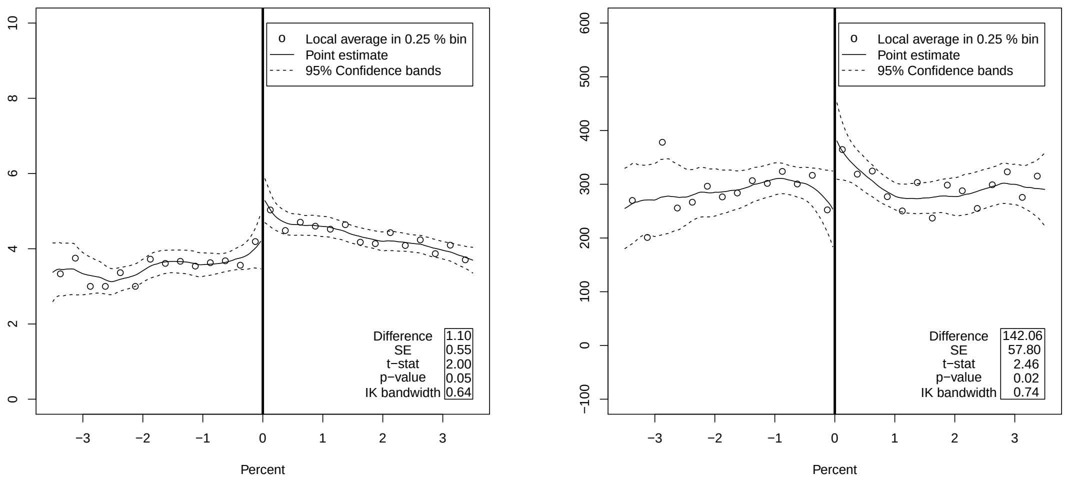

(A) Number of parties winning seats.

(B) Investment per capita in Euro (trimmed).

Consequently, if comparing observations to the left and to the right of the threshold provides me with a good estimate of the average within-observation treatment effect, the estimate of the treatment effect for the number of parties winning seats should be close to 1. The estimate does not mechanically have to equal 1, as there might be slight differences in the average of the baseline number of parties winning seats, i.e., the number of parties without party i, for the compared subgroups on both sides of the threshold. However, an estimate considerably different from 1 would point to a discontinuity in the baseline number of parties winning seats, which in itself would cast doubt on the validity of the identification strategy. The main estimate for the effect is 1.10. With the inclusion of controls or with a variation of the bandwidth, the estimates range from 0.97 to 1.15. These estimates suggest that the empirical strategy passes this first sanity check.

The panel for the number of parties winning seats in Figure 4.2 shows that on both sides of the threshold the number of parties increases with proximity to the threshold. This observation is not surprising, as the party closest to the threshold out of many parties will generally be closer than the closest party out of few. Nevertheless, this fact does not threaten the identification strategy, as the successful sanity check demonstrates.

Figure 4.3: Local linear regression conducted separately on both sides of the threshold. Outcome: a dummy variable indicating whether the party dropped out in $t+1$ and the vote share in the next election (to deal with selection bias, I estimate bounds on the treatment effect by proxying for parties that dropped out in $t+1$ in a different way - for details see Section III.iii). The box in the bottom right corner of each panel contains the estimate of the discontinuity (treatment effect).

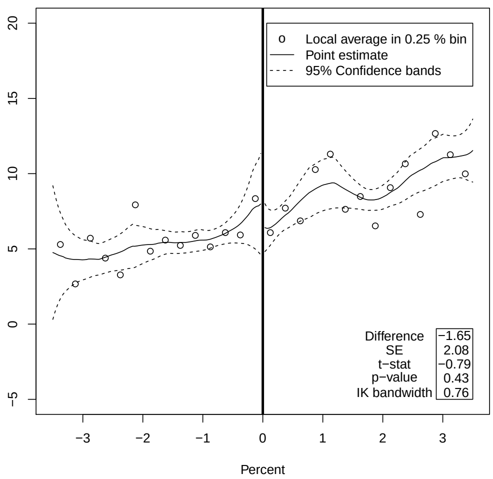

(a) Vote share $t+1$ (lower bound).

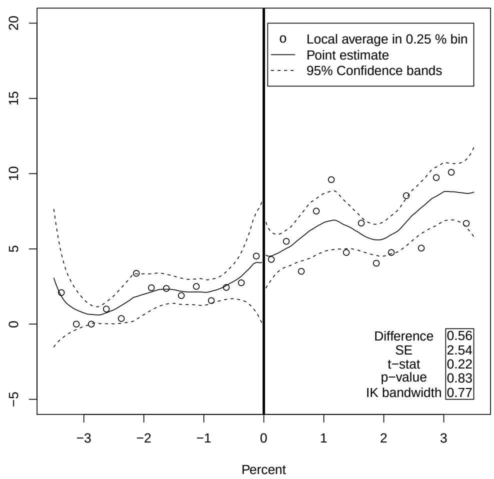

(b) Vote share $t+1$ (upper bound).

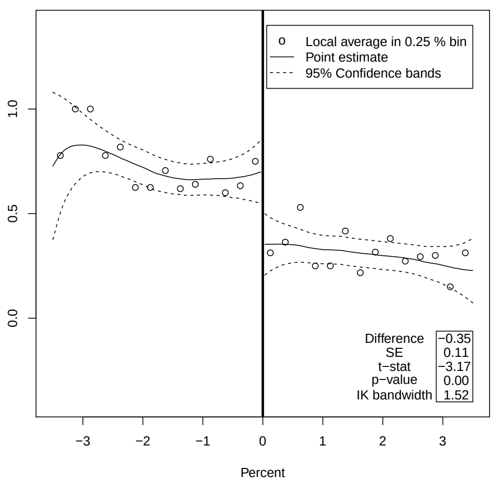

(c) Party not running $t+1$.

I find a considerable effect with regard to investment spending. One party gaining enough votes to be represented in the municipal council, so that overall one additional party is represented, leads to a statistically significant average increase in per capita investment spending of about 142 Euros in the years after the election. This estimate is an average effect. Given discontinuity for the number of parties winning seats displayed in Figure 4.2, the average municipality close to the threshold changes from 4 to 5 parties in the municipal council when gaining an additional party. However, the effect of a third additional party is potentially different from the effect of a sixth additional party. In Section IV.i I investigate the composition of the average effect.

Compared to the average level of per capita investment spending in the trimmed dataset of roughly 295 Euros, this treatment effect represents a sizable increase of almost 50 percent. The estimated effect does not change significantly with the inclusion of controls. The effect does however react to a variation of the bandwidth. It ranges from 201 Euros with half the optimal bandwidth to 102 Euros with double the optimal bandwidth, all the while being significantly different from zero. This variation suggests that, in the data, the difference between parties that won seats and those that did not is more pronounced for parties closer to the threshold.

The graphical representation of the discontinuity in Figure 4.2 also reveals that this result seems to be driven by the observations within a 1-percentage point window around the threshold. Within this window, investment spending is lower for municipalities with parties narrowly missing representation the closer the party missed. The mirroring statement is true for municipalities with parties that narrowly won. Further away from the threshold, there does not seem to be a shift in the level of investment spending depending on treatment status. Put differently, there does not seem to be a clear difference between municipalities that are on different sides of the threshold but sufficiently far away from it.

Table 4.3 contains the estimation results with regard to the question whether treatment status affects the propensity to run in the next election. The dependent variable is a dummy that equals 1 if the party drops out in the next election. I find a sizable and statistically significant effect of -0.35. Winning seats on the municipal council reduces the propensity to drop out in the next election by 35 percent. Again, this effect remains virtually unchanged by the inclusion of population and election round dummies as controls. And again, a variation of the bandwidth does have an impact on the size, though not on the statistical significance of the effect, with estimates ranging from −0.49 for half the optimal bandwidth to −0.33 for double the optimal bandwidth.

With regard to dropping out in the next election, the corresponding panel in Figure 4.3 shows a weak negative relationship between x_i and the propensity to drop out. As a larger value of x_i corresponds to a larger vote share, this relationship is not surprising. A party gaining more votes is more likely to run in the next election. At the threshold, there is a clear shift in the level of parties dropping out. This shift is relatively large, compared to the flat slope of the local linear fit. The comparison of shift and slope suggests that, for a party close to the threshold, in determining whether it should run again there is more weight on whether it was able to win representation compared to the relative success in the election in terms of votes .

For treatment effects on the vote share in the next election, Table 4.4 contains the estimate for the lower bound and Table 4.5 contains the estimate of the upper bound. The baseline point estimate of the lower bound, −1.65, is not statistically significant. Estimates remain not significant with the inclusion of controls or with a variation of the bandwidth. All point estimates remain negative, except the estimate with half the optimal bandwidth. For the upper bound estimate, results are similar. The baseline estimate of 0.56 is not statistically significant, and nor are the estimates with controls and different bandwidths. All point estimates are positive.

The graphs for the lower bound and the upper bound estimation of the vote share incumbency effect in Figure 4.3 show a clear positive relationship of the forcing variable and the vote share in the next election. This result is not surprising, as the forcing variable represents a normalized version of the vote share in the current election, which should be a good predictor of the vote share next election. Both for the lower bound and upper bound estimation, the path of the local linear estimate seems comparatively smooth and continuous at the threshold, compared to the overall variation away from the threshold. Hence, the graphical representation does not suggest the existence of an incumbency advantage.

Table 4.1: Estimating the treatment effect using local linear regression - Outcome: Number of parties winning seats.

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	1.10 **	1.02 *	1.15 **	1.08 *	0.99	0.97 **
	(0.55)	(0.56)	(0.55)	(0.56)	(0.81)	(0.35)
Bandwidth	0.64	0.64	0.64	0.63	0.32	1.28
Within bandwidth						
Obs	143	142	142	142	68	269
Clusters	123	122	122	122	63	210
Overall						
Obs	571	571	571	571	571	571
Clusters	371	371	371	371	371	371
Controls						
Population	No	Yes	No	Yes	No	No
Election round	No	No	Yes	Yes	No	No

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. The treatment effect corresponds to the $\widehat{\tau}$ that solves (4.4). The estimates $\widehat{\alpha}$, $\widehat{\beta}$, $\widehat{\gamma}$, and $\widehat{\delta}$ are not reported. Bandwidth in columns 1-4 is computed according to Imbens and Kalyanaraman (2011). The bandwidth in columns 5 and 6 is half and twice the bandwidth of column 1, respectively.

Table 4.2: Estimating the treatment effect using local linear regression - Outcome: Investment per capita in Euros (trimmed).

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	142.06 **	141.53 **	140.80 **	142.30 **	201.05 **	102.10 **
	(57.80)	(58.76)	(52.97)	(53.66)	(76.00)	(40.09)
Bandwidth	0.74	0.74	0.67	0.67	0.37	1.48
Within bandwidth						
Obs	154	154	145	145	77	292
Clusters	129	129	122	122	71	218
Overall						
Obs	542	542	542	542	542	542
Clusters	358	358	358	358	358	358
Controls						
Population	No	Yes	No	Yes	No	No
Election round	No	No	Yes	Yes	No	No

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. The treatment effect corresponds to the $\widehat{\tau}$ that solves (4.4). The estimates $\widehat{\alpha}$, $\widehat{\beta}$, $\widehat{\gamma}$, and $\widehat{\delta}$ are not reported. Bandwidth in columns 1-4 is computed according to Imbens and Kalyanaraman (2011). The bandwidth in columns 5 and 6 is half and twice the bandwidth of column 1, respectively.

Table 4.3: Estimating the treatment effect using local linear regression - Outcome: Party not running $t+1$ (dummy variable).

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	-0.35 **	-0.33 **	-0.34 **	-0.33 **	-0.49 **	-0.31 **
	(0.11)	(0.11)	(0.11)	(0.11)	(0.16)	(0.08)
Bandwidth	1.52	1.52	1.52	1.49	0.76	3.04
Within bandwidth						
Obs	303	303	303	295	158	494
Clusters	226	226	226	222	132	327
Overall						
Obs	549	549	549	549	549	549
Clusters	353	353	353	353	353	353
Controls						
Population	No	Yes	No	Yes	No	No
Election round	No	No	Yes	Yes	No	No

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. The treatment effect corresponds to the $\widehat{\tau}$ that solves (4.4). The estimates $\widehat{\alpha}$, $\widehat{\beta}$, $\widehat{\gamma}$, and $\widehat{\delta}$ are not reported. Bandwidth in columns 1-4 is computed according to Imbens and Kalyanaraman (2011). The bandwidth in columns 5 and 6 is half and twice the bandwidth of column 1, respectively.

TABLE 4.4: Estimating a lower bound on the treatment effect using local linear regression - Outcome: Vote share $t+1$ (lower bound: parties that dropped out in $t+1$ are assigned value 0).

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	-1.65	-1.48	-1.66	-1.51	0.96	-1.62
	(2.08)	(2.05)	(1.99)	(1.97)	(2.43)	(1.44)
Bandwidth	0.76	0.76	0.76	0.75	0.38	1.52
Within bandwidth						
Obs	158	158	158	158	79	301
Clusters	132	132	132	132	73	225
Overall						
Obs	549	549	549	549	549	549
Clusters	353	353	353	353	353	353
Controls						
Population	No	Yes	No	Yes	No	No
Election round	No	No	Yes	Yes	No	No

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. The treatment effect corresponds to the $\widehat{\tau}$ that solves (4.4). The estimates $\widehat{\alpha}$, $\widehat{\beta}$, $\widehat{\gamma}$, and $\widehat{\delta}$ are not reported. Bandwidth in columns 1-4 is computed according to Imbens and Kalyanaraman (2011). The bandwidth in columns 5 and 6 is half and twice the bandwidth of column 1, respectively.

Table 4.5: Estimating an upper bound on the treatment effect using local linear regression - Outcome: Vote share t+1 (upper bound: parties that dropped out in $t+1$ are assigned vote share t for vote share $t+1$).

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	0.56	0.63	0.49	0.55	2.97	0.11
	(2.54)	(2.52)	(2.44)	(2.44)	(3.12)	(1.73)
Bandwidth	0.77	0.77	0.76	0.76	0.38	1.53
Within bandwidth						
Obs	158	158	158	158	79	303
Clusters	132	132	132	132	73	226
Overall						
Obs	549	549	549	549	549	549
Clusters	353	353	353	353	353	353
Controls						
Population	No	Yes	No	Yes	No	No
Election round	No	No	Yes	Yes	No	No

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. The treatment effect corresponds to the $\widehat{\tau}$ that solves (4.4). The estimates $\widehat{\alpha}$, $\widehat{\beta}$, $\widehat{\gamma}$, and $\widehat{\delta}$ are not reported. Bandwidth in columns 1-4 is computed according to Imbens and Kalyanaraman (2011). The bandwidth in columns 5 and 6 is half and twice the bandwidth of column 1, respectively.

Determinants of the average treatment effect

As previously discussed, my main set of estimation results for the treatment effects represent average effects. However, the effect of a party narrowly making it onto the municipal council might differ depending on whether the change in treatment status changed the number of parties in the municipal council from 2 to 3 parties or from 6 to 7 parties. Similarly, there might be a different effect depending on whether an additional party considerably changed majorities and therefore potential coalitions in the municipal council. In this section, I try to disentangle the estimated average treatment effects and look into potential channels behind my results.

Technically, I am interested in how the treatment effect varies conditional on some covariate, such as, for example, the number of parties. To investigate this question, I estimate a specification in which z_i, the covariate in question, is also interacted with the switch in treatment status at the threshold,

$$\min_{\alpha,\tau,\beta,\beta_2,\gamma,\gamma_2} \sum_{\{i:-h\leq x_i\leq h\}} (Y_i-\alpha-\tau D_i-\beta x_i-\gamma x_i D_i-\beta_2 z_i-\gamma_2 z_i D_i)^2 K(x_i,0,h). \quad (4.6)$$

The crucial difference to specification (4.5) is the interaction term $z_i D_i$, which allows me to estimate the treatment effect conditional on z_i:

$$Y_i(D_i=1|z_i) - Y_i(D_i=1|z_i) = \tau + \gamma_2 z_i.$$

This specification means that the estimated treatment effect varies linearly in z_i by construction.

To obtain a heterogeneous treatment effect by the number of parties winning seats I use the baseline number of parties with seats, which is the number of parties winning seats without counting the party in question. If party i enters the municipal council as the fifth party winning seats, then the baseline number of parties is 4 and the treatment effect for party i is $\tau + \gamma_2 \cdot 4$. I would expect treatment effects to be stronger with a lower baseline number of parties, as treatment comes with larger change to the number of parties winning seats in relative terms, giving more bargaining power to a party narrowly entering the municipal council. For investment per capita, the coefficient of the interaction effect is -37, which therefore has the expected sign, and is statistically significant. According to the coefficient, one additional baseline party would imply a decrease of the treatment effect of 37

Euros, which however remains positive in all cases but the extreme case with 9 baseline parties. In contrast, with regard to how treatment affects the propensity to drop out in the next election, I cannot reject the null hypothesis that the effect does not vary by the baseline number of parties. Apart from not being statistically significant, the estimate for dropping out in the next election is 0. The treatment effects by baseline number of parties are depicted in the left panels in Figure 4.3.

The number of baseline parties can be considered a measure of how an additional party affects or would affect the political setup on the municipal council. Next, I use the Herfindahl index as an alternative measure of the political setup. The Herfindahl index is a measure of concentration. In the context of Economics it is more commonly used in the field of Industrial Organization to measure market concentration. In Political Science and Political Economy it is used to measure the concentration or fractionalization of a legislature. To capture the concentration in a municipal council I calculate

$$(\text{Herfindahl index})_M := \sum_{i \in M} \left(\frac{s_i}{\mathbf{s}_M} \right)^2 \in (0, 1], \tag{4.7}$$

where s_i represents the seats obtained by party i in municipality M and $\mathbf{s}_M$ represents the total number of seats in the municipal council.

By measuring fractionalization, the Herfindahl index captures the extent of different interests represented in the municipal council. A lower Herfindahl index indicates a more fractionalized municipal council, where the maximal Herfindahl index of 1 indicates that the council is not fractionalized and all seats are held by a single party. Achieving a majority in a vote might be more difficult in a more fractionalized municipal council and in which compromise might be more important. However, the Herfindahl index does not account for variations in the political compatibility between parties. While an ideal measure of the political setup in the municipal council would account for political compatibility, I do not think a sensible measure can be constructed for the municipal election context with the available data.

Calculating the counterfactual allocation of seats, I can measure how much an individual party list affected the setup on the municipal council. Much in the same way as for the number of parties winning seats, I can observe $Y_i(D_i = 1) - Y_i(D_i = 0)$ for an individual party i for the Herfindahl index. With Y_i

FIGURE 4.3: Differential treatment effects obtained by interacting the treatment dummy variable with the baseline number of parties winning seats (without counting the party in question) or with Δ Herfindahl index, the change in the Herfindahl index attributable to the party in question. Dots/the solid line represent/s the point estimate of the treatment effect. Bars/the dashed line represent/s the 95 percent confidence interval.

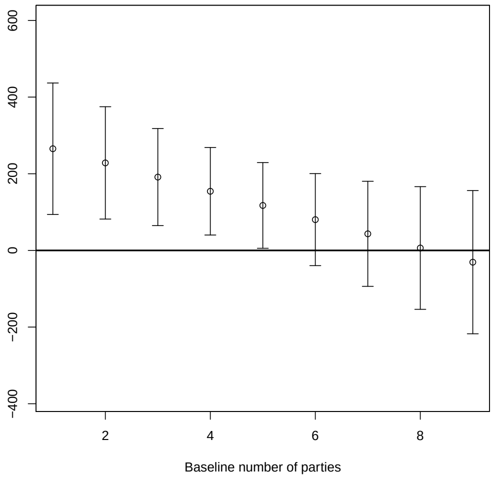

(A) Effect on (trimmed) Investment per capita in Euros by baseline number of parties.

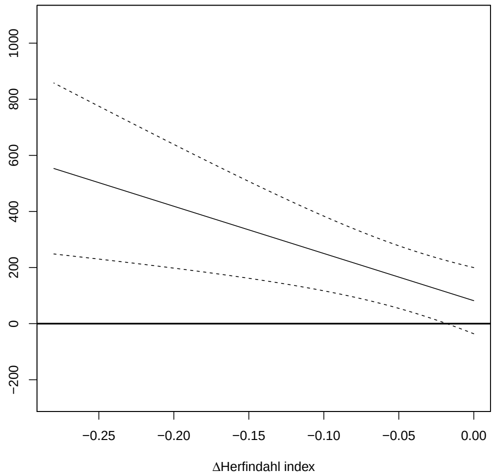

(B) Effect on (trimmed) Investment per capita in Euros by Δ Herfindahl index.

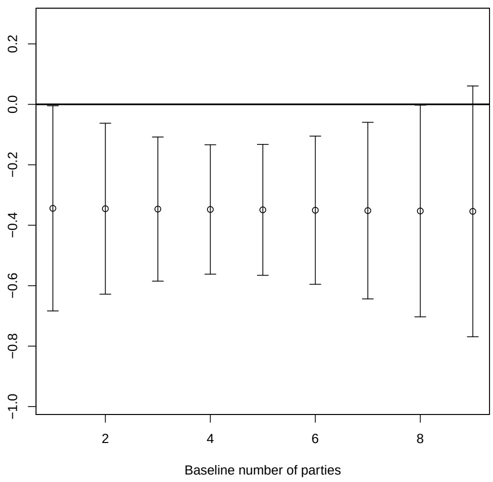

(C) Effect on not running in the next election by baseline number of parties.

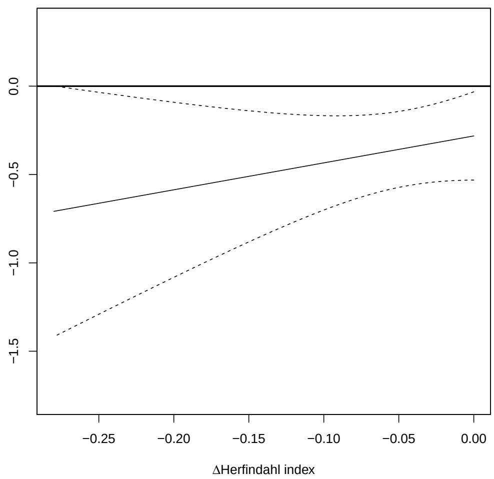

(D) Effect on not running in the next election by Δ Herfindahl index.

Table 4.6: Estimating differential treatment effects using local linear regression - Outcome: Investment per capita in Euros (trimmed). The baseline number of parties winning seats (without the party in question) and Δ Herfindahl index (the change in the Herfindahl index attributable to the party in question) are included with and without an interaction effect.

	(1)	(2)	(3)	(4)
Treatment effect	142.39 **	302.34 **	141.60 **	81.88
	(58.22)	(101.72)	(58.18)	(60.19)
Baseline n. of parties w. s.	3.91	23.85 *		
	(8.79)	(12.32)		
Baseline n. of parties w. s. $\times D_i$		-37.00 **		
		(17.88)		
Δ Herf. ind.			82.07	1007.55 **
			(375.55)	(287.30)
Δ Herf. ind. $\times D_i$				-1683.25 **
				(593.89)
Bandwidth	0.74	0.74	0.74	0.74
Within bandwidth				
Obs	154	154	154	154
Clusters	129	129	129	129
Overall				
Obs	542	542	542	542
Clusters	358	358	358	358

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. Estimates in columns 1 and 3 are the solutions to (4.5). Estimates in columns 2 and 4 are the solutions to (4.6). The treatment effect corresponds to $\widehat{\tau}$. The estimates $\widehat{\alpha}$, $\widehat{\beta}$, and $\widehat{\gamma}$ are not reported. Bandwidth is computed according to Imbens and Kalyanaraman (2011).

standing for the Herfindahl index, the variable "Δ Herfindahl index" measures the individual treatment effect $Y_i(D_i = 1) - Y_i(D_i = 0)$.

With respect to the interaction effect of treatment and Δ Herfindahl index I would expect a larger effect in absolute terms for larger absolute values of Δ Herfindahl index. Put differently, I would expect parties which, by narrowly making it onto the municipal council, cause a larger change in its political setup to also cause a larger change in investment spending in that municipality. Similarly, I would

Table 4.7: Estimating differential treatment effects using local linear regression - Outcome: Not running in $t+1$ (dummy variable). The baseline number of parties winning seats (without the party in question) and Δ Herfindahl index (the change in the Herfindahl index attributable to the party in question) are included with and without an interaction effect.

	(1)	(2)	(3)	(4)
Treatment effect	-0.35 **	-0.34 *	-0.35 **	-0.28 **
	(0.11)	(0.21)	(0.11)	(0.13)
Baseline n. of parties w. s.	-0.01	-0.01		
	(0.03)	(0.03)		
Baseline n. of parties w. s. $\times D_i$		0.00		
		(0.04)		
Δ Herf. ind.			-0.42	-1.27
			(0.75)	(1.02)
Δ Herf. ind. $\times D_i$				1.52
				(1.47)
Bandwidth	1.52	1.52	1.52	1.52
Within bandwidth				
Obs	303	303	303	303
Clusters	226	226	226	226
Overall				
Obs	549	549	549	549
Clusters	353	353	353	353

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. Estimates in columns 1 and 3 are the solutions to (4.5). Estimates in columns 2 and 4 are the solutions to (4.6). The treatment effect corresponds to $\widehat{\tau}$. The estimates $\widehat{\alpha}$, $\widehat{\beta}$, and $\widehat{\gamma}$ are not reported. Bandwidth is computed according to Imbens and Kalyanaraman (2011).

expect the probability that those parties drop out in the next election to be more affected by treatment.

For the interaction effect of treatment and Δ Herfindahl index with respect to investment per capita I find a statistically significant coefficient with the expected sign. For observations with the average value of Δ Herfindahl index -0.06, i.e., for observations which by making it onto the council lower the Herfindahl index by 0.06, the treatment effect is roughly 183 Euros. The coefficient of the interaction effect with respect to dropping out next election is not statistically significant. It does have the expected sign, however, which would imply that the treatment effect for the propensity to drop out is larger in absolute terms if treatment status leads to a larger change in the Herfindahl index. The treatment effects by Δ Herfindahl index are depicted in the right panels of Figure 4.3.

Robustness checks

As pointed out by Caughey and Sekhon (2011) and Eggers, Fowler, Hainmueller, Hall, and Snyder (2015), among others, it is important to test the validity of the identifying assumption when applying the RD methodology to an election setting. One test they propose is to check for a discontinuity in the density of the sorting variable at the threshold. Such a discontinuity would cast doubt on the identifying assumption.

I test for a discontinuity in the density using methods proposed by McCrary (2008). The idea of this test is to test for a discontinuity in the histogram of the forcing variable. One bin corresponds to an observation, the value of the forcing variable is given by the bin center and the outcome variable is given by the number of observations within the bin. As in the baseline specification for my main estimation results, the test uses local linear regression with a triangular kernel. McCrary (2008) comments on bandwidth selection but his test does not require a specially tailored bandwidth. Again, as in my baseline specification, I use the bandwidth selection method of Imbens and Kalyanaraman (2012).

The results of the test are reported in Table 4.9. The test does not reject the null hypothesis, no discontinuity in the density, for any of the three test bin sizes, and with either the full dataset or the trimmed dataset. A graphical representation of the test for bins with a 0.1 percentage point size can be found in Figure 4.3, the p-values of the tests of Table 4.9 can also be found in Figure 4.4.

Another important test of the identifying assumption is to check for discontinuities in pre-treatment covariates. Covariates that should not be affected by the treatment and exhibit a discontinuity at the threshold would indicate that observations close to the threshold but on different sides are not comparable, contrary to the identifying assumption.

I test for a discontinuity in election variables, such as the number of people eligible to vote, the total number of seats in the municipal council, whether the party list in question belongs to a major political party in Germany, and the population of the municipality.[9] The results of these tests are presented in Figure 4.4. The panel on the left contains the regression for turnout in the election, as an example. The panel on the right contains the p-values for the examined covariates. The null hypothesis, no discontinuity at the threshold, is not rejected for any of these covariates.

Ideally, I would also test for discontinuities in lags of the main outcome variables. In other words, I would test for discontinuities in whether a party was running in the previous election and in investment spending in the years prior to the election. Unfortunately, I cannot conduct these tests, as the necessary data is not available for a large part of my sample. This part includes the 1994 election round, as there was no municipal election of this kind prior to 1994 and as the data for municipal finances is only available from 1995 onwards. Similarly, it also includes municipalities that underwent reorganization in the years prior to the election.

An important robustness check when using local linear regression is to investigate the sensitivity of the results with respect to the choice of the bandwidth. I use the optimal bandwidth proposed by Imbens and Kalyanaraman (2012). My main set of estimation results in Tables 4.2 and 4.3 also contain results for half and twice the optimal bandwidth. The sensitivity of the estimate with respect to the bandwidth is depicted in Figure 4.5. Point estimates of the treatment effect are smaller for larger bandwidths but remain statistically significant for all bandwidths in that range, suggesting that my results are not critically driven by the bandwidth choice.

When using an RD approach, a common robustness check is to estimate the model at pseudo-thresholds, at which treatment status does not change, as it does at the true threshold. If there is a significant effect at many pseudo-thresholds, it is questionable whether a significant effect at the true threshold can be attributed to the change in treatment status. Including the estimation at the true threshold 0, I estimate the baseline model for 21 different thresholds, going from -2 to 2 in 0.1 increments. Figure 4.6 contains the z-statistics for these regressions. For investment spending as the outcome, I find one statistically significant discontinuity at 0.9. For not running in the next election as the outcome, I find one statistically significant discontinuity at the 10 percent level at 1.7. However, under the null

[9] As major parties I classify the previously mentioned Christian Democratic “CDU”, the Social Democratic “SPD”, the Greens, the pro-business “FDP”, and the far-left successor of the ruling party of the GDR: “PDS”. Where a major political party merged with another party in a municipality, the resulting joint party list is not classified as major.

hypothesis, in which the z-statistic is distributed according to a standard normal distribution, i.e., with mean 0, finding one significant value in 20 regressions is expected.

Interestingly, the estimate of the discontinuity is not significant for pseudo-thresholds close to the true threshold, i.e., for -0.1 and 0.1. This result indicates that there are observations[10] both very close to the threshold and with a substantial difference in the average outcome variable contributing to the results of the baseline estimation.

Instead of local linear regression, many researchers use polynomial regression in RD frameworks. For robustness, I also estimate a polynomial regression. The results can be found in Tables 4.10 and 4.11. For polynomials of order 3 to order 6 I find statistically significant effects, at least at the 10 percent level, that are comparable to the estimates in the local linear regression in terms of magnitude.

In Section III.iii I present the reasons for choosing investment spending rather than overall spending as my measure of public spending. Nevertheless, I should find at least comparable effects for overall spending, as investment spending contributes to overall spending. In Table 4.12 I report estimation results with (trimmed) overall gross spending as the outcome variable instead of investment spending. For the baseline specification I do not find a statistically significant treatment effect. The point estimate of 114 Euros is, however, comparatively close to the point estimate of 142 Euros for investment spending. The fact that the point estimate is slightly lower might be caused by the somewhat larger bandwidth selected for overall spending. At half the optimal bandwidth, the treatment effect is significant at the 10 percent level and similar in size to the treatment effect for investment spending. It is considerably smaller and not significant at twice the optimal bandwidth. With respect to controls, I find that the baseline estimate for overall spending is only limitedly robust. Unlike the estimate for investment spending, the estimate for overall spending reacts much more strongly to the inclusion of population size and election year dummies. All in all, keeping in mind the link between the point estimate and the bandwidth, I find similar effects for overall spending and investment spending. This result suggests that the overall spending effect reflects the investment spending effect, though with added noise, which provides some justification for the choice of investment spending as my measure of public spending.

[10] There are 22 observations within $x_i \in [-0.1, 0.1]$.

Sensitivity to trimming

As mentioned in Section III.v, for the estimation with investment per capita as the outcome variable I trim the data at the 95th percentile. This procedure might cause some concern that the estimated effects are biased or that I only find significant effects because, through trimming, I investigate a sample in which finding an effect is more likely. This phenomenon is referred to as “selection on the dependent variable”. I need to be careful in this regard, as the number of parties in the municipal council might be more likely to affect the small and medium level type of investment spending described in Section III.ii than very large investment projects.

Graphically, the distribution of investment per capita and some robustness checks with regard to trimming are presented in Figure 4.7. As depicted in the upper left panel, estimating the baseline specification with the untrimmed dataset leads to an estimate that is not too far in magnitude from the main estimation result for the treatment effect of 142 Euros, though it is not statistically significant. The distribution of the untrimmed investment per capita is shown in the top right panel of the figure. The 95th percentile roughly coincides with Tukey’s upper fence, i.e., the set of outliers in Tukey’s sense roughly coincides with the top 5 percent of the data. The bottom left panel shows how the point estimate changes if a percentile other than the 95th is chosen for trimming, ranging from trimming at the 90th percentile to no trimming at 100 percent. Other than when the data is not trimmed or trimmed at a percentile higher than the 99th, the estimate of the treatment effect is relatively stable in size and significant. The bottom right panel contains an estimation of the discontinuity at the threshold with a dummy, whether the observation was kept or trimmed. The null hypothesis, that the data is not discontinuously trimmed at the threshold, is not rejected.

An alternative to trimming is to use investment in logs. The results of this estimation are reported in Table 4.13. The estimates of the baseline specification and the specifications with population as control are positive and significant at the 10 percent level. However, the estimates are no longer significant with the inclusion of controls and at half the optimal bandwidth. While not very robust, these results at least give some indication that the estimates obtained for the trimmed dataset are not just spurious, and are not due to trimming.

Nevertheless, it is important to remember that the data is trimmed and represents the bottom 95 percent of the sample when interpreting the main result. Consequently, for my main result, I find an average effect of 142 Euros of additional investment spending with one additional party in the municipal council for municipalities that do not have a very large amount of investment spending.

Conclusion

Recently, many researchers have used close elections as natural experiments. In majoritarian election systems with two parties, elections with an almost 50-50 outcome are used to identify the causal effect of incumbency. In proportional representation electoral systems, discontinuities in the seat allocation method can also be used for identification. In a proportional representation context, there are two margins of incumbency: the intensive margin, capturing how many seats a party won, and the extensive margin, capturing whether the party won seats at all. While intensive margin effects have been studied to a wide extent, there is little evidence with regard to extensive margin effects. My paper is the first study focused on the causal estimation of the effect of a small party winning representation in a proportional representation system. I apply a regression discontinuity (RD) approach to municipal level data from the German state of Thuringia.

On the municipal level, I estimate the effect it has on public spending if a small party narrowly wins seats and thereby increases the overall number of parties with seats on the municipal council. The Political Economy literature generally predicts a positive relationship of legislative fractionalization and spending and therefore a positive effect. Using (trimmed) investment spending as my measure of public spending, I find a large and statistically significant positive effect. One additional party increases investment spending by 50 percent of the average level of spending in my dataset. Decomposing this average effect, I find that the effect is stronger where a party winning seats leads to a larger change to the composition of the municipal council. A graphical inspection of the discontinuity reveals that the estimated effect is driven by observations within a relatively narrow window around the threshold. Therefore, I believe that further research to investigate this relationship is warranted. Nevertheless, a large variety of robustness checks and the fact that the decomposition of the effect produces results in line with theoretical predictions give me confidence that the estimated effect is not just a statistical anomaly.

On the party level, I find that parties that narrowly won representation are 35 percent less likely to drop out in the next election. Similar effects have been observed in RD studies for individual candidates in the majoritarian context but not for parties in the proportional representation context. This discontinuity leads to a sample selection problem that precludes me from obtaining an unbiased estimate of a potential vote share incumbency effect. Conducting a bounding exercise instead, I find no evidence of an incumbency effect on the vote share in the next election. The lower bound point estimate is negative, the upper bound point estimate is positive, and neither is statistically significant. This result is in contrast to Liang (2013), who finds an incumbency advantage of approximately 0.7 percent at the extensive margin in the proportional representation context. However, given the magnitude of his estimate, it could be that I simply lack statistical power to find a similar sized statistically significant effect.

Overall, I find sizable effects at the extensive margin. This finding suggests that a small shift in the election outcome reflecting only a slight shift in the preferences of the electorate can have a comparatively large impact, if it leads to an additional party winning seats. A better understanding of these effects can be useful in the design of political institutions.

Table 4.8: Summary statistics.

	Mean	Median	SD	Min	Max	Obs
Forcing variable - distance to the threshold in Percent: x_i	0.44	0.41	1.77	-3.65	3.67	571
Number of parties winning seats	4.04	4.00	1.29	1.00	10.00	571
Baseline number of parties winning seats	3.45	3.00	1.23	1.00	9.00	571
Number of parties running	4.74	5.00	1.51	2.00	11.00	571
Number of parties running $t+1$	4.23	4.00	1.50	2.00	11.00	549
Herfindahl index	0.38	0.34	0.14	0.12	1.00	571
Δ Herfindahl index	-0.06	-0.04	0.05	-0.28	0.00	571
Major political party	0.38	0.00	0.48	0.00	1.00	571
Party dropped out in $t+1$ (dummy variable)	0.47	0.00	0.50	0.00	1.00	549
Party merged with other party in $t+1$ (dummy variable)	0.03	0.00	0.18	0.00	1.00	549
Party split up in $t+1$ (dummy variable)	0.01	0.00	0.07	0.00	1.00	549
Population (election year)	7061.00	2614.00	17913.85	125.00	202400.00	571
Number of seats on council	15.25	14.00	8.03	6.00	50.00	571
Eligible to vote	5676.00	2083.00	14402.43	101.00	164800.00	571
Voters	3210.00	1364.00	7321.18	90.00	79970.00	571
Turnout	68.23	69.45	12.06	38.32	100.00	571
Turnout $t+1$	60.60	60.41	10.52	38.32	98.13	549
Percentage of invalid votes	4.50	4.06	2.05	0.00	20.66	571
Percentage of invalid votes $t+1$	4.45	3.89	2.21	0.00	20.66	549
Investment per capita in Euros (trimmed)	295.40	271.10	157.44	19.26	752.10	542
Investment per capita in Euros (untrimmed)	348.60	277.40	381.25	19.26	7092.00	571
Log Investment per capita in Euros	5.60	5.63	0.70	2.96	8.87	571
Gross spending per capita in Euros (trimmed)	1187.00	1141.00	296.21	550.50	2058.00	542
Gross spending per capita in Euros (untrimmed)	1301.00	1162.00	727.14	550.50	9600.00	571

Table 4.9: Histogram test, as proposed by McCrary (2008). Columns 1-3 contain tests of the full dataset with various bin sizes. Columns 4-6 contain tests of the data trimmed with respect to investment per capita.

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	-0.18	0.64	0.66	0.35	0.88	1.22
	(0.47)	(1.88)	(1.90)	(0.63)	(1.75)	(1.77)
Bin size	0.025	0.050	0.100	0.025	0.050	0.100
Bandwidth	2.82	1.04	1.04	1.47	1.03	1.03
Within bandwidth						
Obs	226	42	20	118	42	20
Overall						
Obs	320	160	80	320	160	80

Significance levels: * 0.1, ** 0.05. Standard errors are robust. Estimates are obtained using local linear regression with a triangular kernel. The treatment effect corresponds to a potential discontinuity in the density of the sorting variable at the threshold. The data for the test is based on a histogram of the forcing variable on the support [-4,4]. Each observation corresponds to a bin. The value of the outcome variable is the number of observations within a bin, the value of the forcing variable is the bin center. Bandwidth is computed according to Imbens and Kalyanaraman (2011).

Table 4.10: Estimating the treatment effect using polynomial regression instead of local linear regression - Outcome: Investment per capita in Euros (trimmed).

	(1)	(2)	(3)	(4)
Treatment effect	125.12 **	161.08 **	166.79 **	196.91 **
	(49.14)	(60.66)	(72.11)	(85.40)
Degree of polynomial	3	4	5	6
Obs	542	542	542	542
Clusters	358	358	358	358
AIC	7026.451	7029.470	7033.453	7036.709

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality.

Table 4.11: Estimating the treatment effect using polynomial regression instead of local linear regression - Outcome: Party not running in $t + 1$ (dummy variable).

	(1)	(2)	(3)	(4)
Treatment effect	-0.43 **	-0.31 *	-0.54 **	-0.66 **
	(0.14)	(0.17)	(0.21)	(0.25)
Degree of polynomial	3	4	5	6
Obs	549	549	549	549
Clusters	353	353	353	353
AIC	720.8352	723.0383	723.5698	726.6073

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality.

Table 4.12: Estimating the treatment effect using local linear regression - Outcome: Gross spending per capita in Euros (trimmed).

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	114.42	90.93	94.79	73.04	176.04 *	49.64
	(79.29)	(77.78)	(76.77)	(76.32)	(93.87)	(61.56)
Bandwidth	1.09	1.03	1.08	1.00	0.54	2.17
Within bandwidth						
Obs	220	208	220	201	120	391
Clusters	179	171	179	165	108	284
Overall						
Obs	542	542	542	542	542	542
Clusters	360	360	360	360	360	360
Controls						
Population	No	Yes	No	Yes	No	No
Election round	No	No	Yes	Yes	No	No

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. The treatment effect corresponds to the $\widehat{\tau}$ that solves (4.4). The estimates $\widehat{\alpha}$, $\widehat{\beta}$, $\widehat{\gamma}$, and $\widehat{\delta}$ are not reported. Bandwidth in columns 1-4 is computed according to Imbens and Kalyanaraman (2011). The bandwidth in columns 5 and 6 is half and twice the bandwidth of column 1, respectively.

Table 4.13: Estimating the treatment effect using local linear regression - Outcome: Log Investment per capita in Euros.

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	0.38 *	0.37	0.32	0.32	0.52	0.27 *
	(0.23)	(0.23)	(0.21)	(0.22)	(0.32)	(0.16)
Bandwidth	0.89	0.89	0.90	0.90	0.45	1.79
Within bandwidth						
Obs	187	187	187	187	101	355
Clusters	156	156	156	156	94	263
Overall						
Obs	571	571	571	571	571	571
Clusters	371	371	371	371	371	371
Controls						
Population	No	Yes	No	Yes	No	No
Election round	No	No	Yes	Yes	No	No

Significance levels: * 0.1, ** 0.05. Standard errors are clustered by municipality. Estimates are obtained using local linear regression with a triangular kernel. The treatment effect corresponds to the $\widehat{\tau}$ that solves (4.4). The estimates $\widehat{\alpha}$, $\widehat{\beta}$, $\widehat{\gamma}$, and $\widehat{\delta}$ are not reported. Bandwidth in columns 1-4 is computed according to Imbens and Kalyanaraman (2011). The bandwidth in columns 5 and 6 is half and twice the bandwidth of column 1, respectively.

Figure 4.3: Robustness check - checking for a discontinuity in the density of the forcing variable as proposed by McCrary (2008). Panels contain a histogram of the forcing variable overlaid with a local linear regression on both sides of the threshold. Estimates of the discontinuity are also reported in Table 4.9.

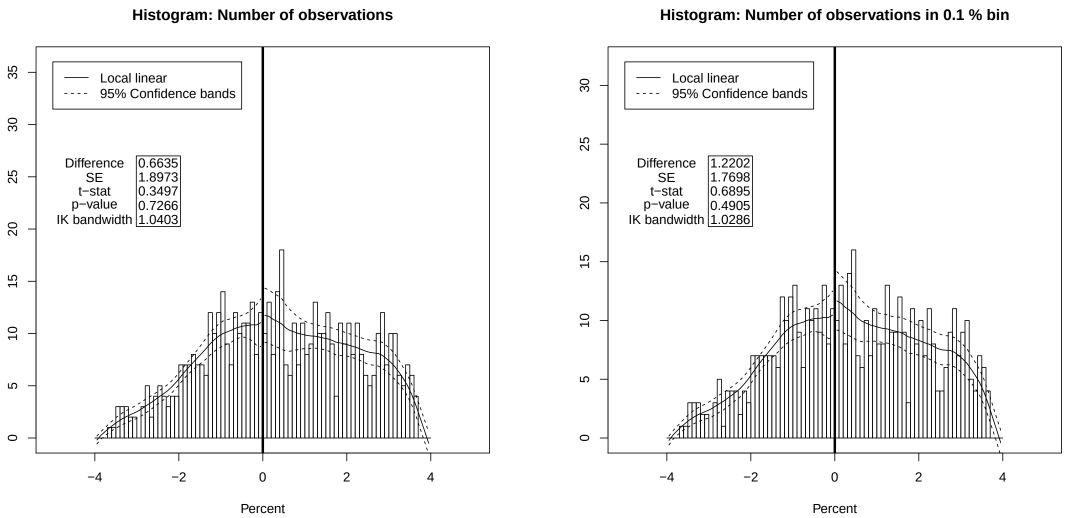

(a) Bin size 0.1 - full dataset.

(b) Bin size 0.1 - data trimmed with respect to investment per capita.

FIGURE 4.4: Robustness check - checking for discontinuous sorting at the threshold.

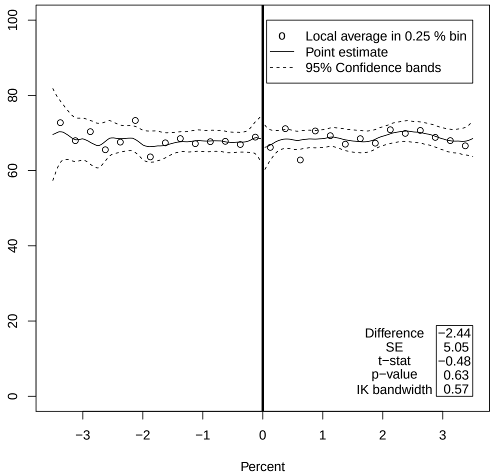

(A) One example of a variable that should not vary discontinuously (i.e., exhibit a significant treatment effect) at the threshold: Turnout $t+1$.

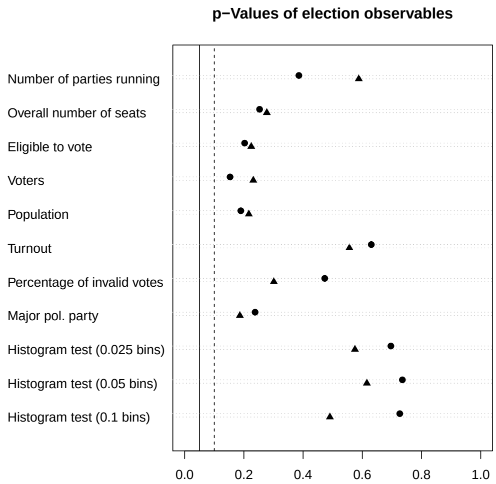

(B) p-values of the estimated discontinuity at the threshold for some election variables that should not exhibit a significant discontinuity at the threshold. Circles represent the full dataset, triangles represent the dataset trimmed with respect to investment per capita. The last three rows contain tests for discontinuities in the density, as in Figure 4.3, for different bin sizes.

Figure 4.5: Robustness check - estimated treatment effect for different bandwidths.

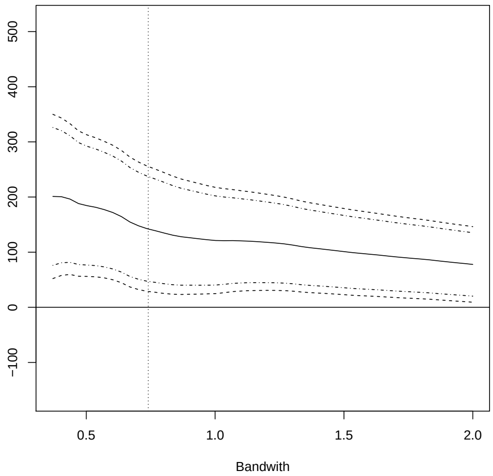

(A) Investment per capita in Euros (trimmed).

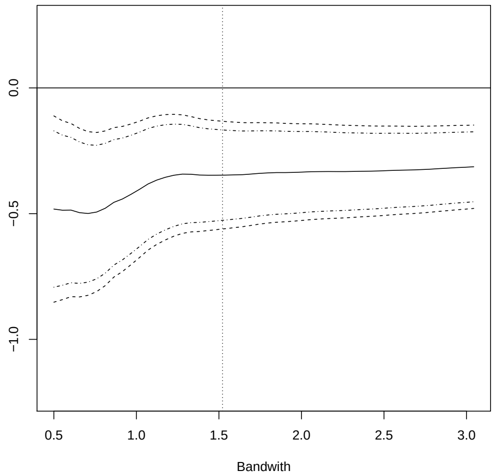

(B) Party not running $t+1$ (dummy variable).

Figure 4.6: Robustness check - checking for significant discontinuities at other values of x_i besides the true threshold. z-value on the vertical axis. Dashed lines mark boundary of 10 and 5 percent significance.

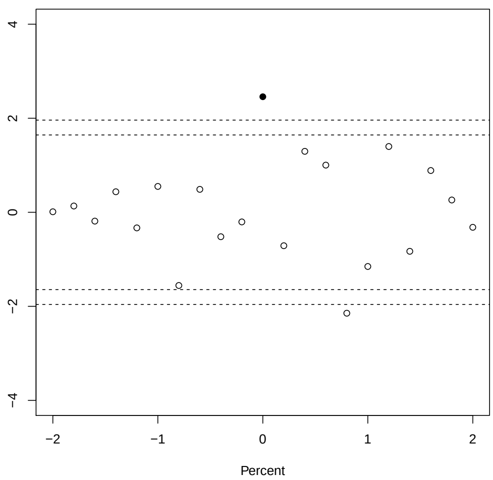

(A) Investment per capita in Euros (trimmed).

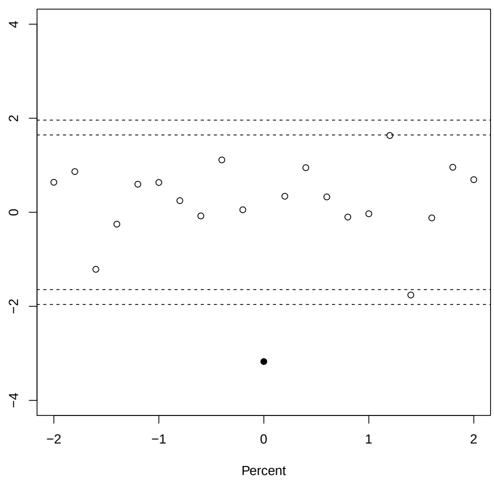

(B) Party not running $t+1$ (dummy variable).

FIGURE 4.7: Robustness check - sensitivity of the estimation results for investment per capita with respect to trimming.

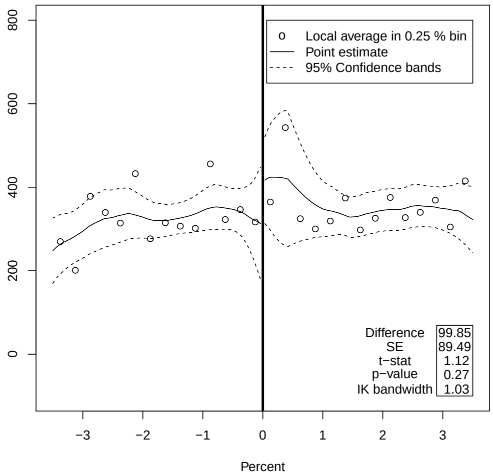

(A) Local linear regression with (untrimmed) investment per capita as outcome variable.

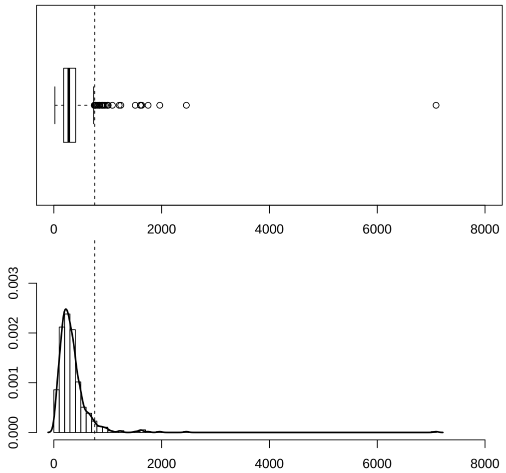

(B) Graphical representation of the outliers for investment per capita. Boxplot at the top. Histogram overlayed with a density plot at the bottom. The dashed line represents the 95th percentile.

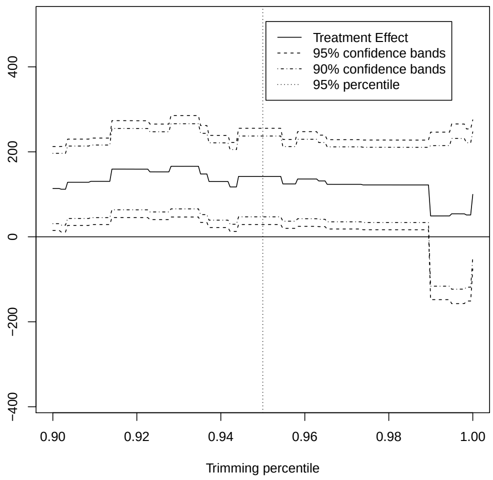

(C) Estimated treatment effect with data trimmed at different percentiles.

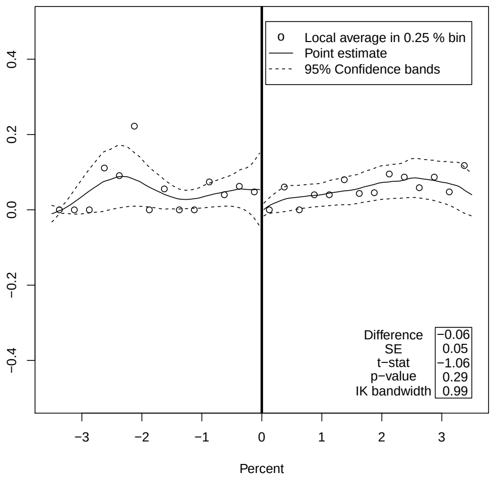

(D) Investigating whether data is discontinuously trimmed at the threshold. Local linear regression with a trimming dummy as the outcome variable.

Bibliography

Aggarwal, G., and J. D. Hartline (2006): "Knapsack auctions," in *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp. 1083–1092. Society for Industrial and Applied Mathematics.

Albouy, D. (2013): "Partisan representation in Congress and the geographic distribution of Federal funds," *Review of Economics and Statistics*, 95(1), 127–141.

Alford, J. R., and D. W. Brady (1993): "Personal and partisan advantage in US congressional elections," in *Congress Reconsidered*, ed. by L. C. Dodd, and B. I. Oppenheimer. CQ Press, Washington DC, 5 edn.

Anari, N., G. Goel, and A. Nikzad (2014): "Mechanism Design for Crowdsourcing: An Optimal 1-1/e Competitive Budget-Feasible Mechanism for Large Markets," in *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pp. 266–275. IEEE.

Austen-Smith, D. (1990): "Information transmission in debate," *American Journal of Political Science*, pp. 124–152.

Austen-Smith, D., and T. J. Feddersen (2009): "Information aggregation and communication in committees," *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1518), 763–769.

Baron, D. P., and R. Myerson (1982): "Regulating a monopolist with unknown costs," *Econometrica*, pp. 911–930.

Bei, X., N. Chen, N. Gravin, and P. Lu (2012): "Budget feasible mechanism design: from prior-free to bayesian," in *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 449–458. ACM.

Börgers, T., and P. Norman (2009): "A note on budget balance under interim participation constraints: the case of independent types," *Economic Theory*, 39(3), 477–489.

Caillaud, B., and J. Robert (2005): "Implementation of the revenue-maximizing auction by an ignorant seller," *Review of Economic Design*, 9(2), 127–143.

Carlsson, M., G. B. Dahl, and D.-O. Rooth (2015): "Do Politicians Change Public Attitudes?," Discussion paper, National Bureau of Economic Research.

Caughey, D., and J. S. Sekhon (2011): "Elections and the regression discontinuity design: Lessons from close us house races, 1942–2008," *Political Analysis*, 19(4), 385–408.

Che, Y.-K., and I. Gale (1998): "Standard auctions with financially constrained bidders," *The Review of Economic Studies*, 65(1), 1–21.

Chen, N., N. Gravin, and P. Lu (2011): "On the approximability of budget feasible mechanisms," in *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pp. 685–699. SIAM.

Chung, K.-S., and J. C. Ely (2002a): "Ex-post incentive compatible mechanism design," *Mimeo, Northwestern University.*

——— (2002b): "Mechanism Design with Budget Constraint," *Mimeo, Northwestern University.*

Coughlan, P. J. (2000): "In defense of unanimous jury verdicts: Mistrials, communication, and strategic voting," *American Political Science Review*, 94(02), 375–393.

Crawford, V. P., and J. Sobel (1982): "Strategic information transmission," *Econometrica*, pp. 1431–1451.

Dizdar, D., A. Gershkov, and B. Moldovanu (2011): "Revenue maximization in the dynamic knapsack problem," *Theoretical Economics*, 6(2), 157–184.

Dobzinski, S., C. H. Papadimitriou, and Y. Singer (2011): "Mechanisms for complement-free procurement," in *Proceedings of the 12th ACM conference on Electronic commerce*, pp. 273–282. ACM.

Doraszelski, U., D. Gerardi, and F. Squintani (2003): "Communication and voting with double-sided information," *Contributions in Theoretical Economics*, 3(1).

Dütting, P., V. Gkatzelis, and T. Roughgarden (2014): "The performance of deferred-acceptance auctions," in *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 187–204. ACM.

Eggers, A. C., A. Fowler, J. Hainmueller, A. B. Hall, and J. M. Snyder (2015): "On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races," *American Journal of Political Science*, 59(1), 259–274.

Ensthaler, L., and T. Giebe (2014a): "Bayesian optimal knapsack procurement," *European Journal of Operational Research*, 234(3), 774–779.

——— (2014b): "A dynamic auction for multi-object procurement under a hard budget constraint," *Research Policy*, 43(1), 179–189.

Farrell, J. (1993): "Meaning and credibility in cheap-talk games," *Games and Economic Behavior*, 5(4), 514–531.

Feddersen, T., and W. Pesendorfer (1998): "Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting," *American Political Science Review*, 92(01), 23–35.

Ferreira, F., and J. Gyourko (2009): "Do Political Parties Matter? Evidence from US Cities," *The Quarterly Journal of Economics*, 124(1), 399–422.

Folke, O. (2014): "Shades of brown and green: party effects in proportional election systems," *Journal of the European Economic Association*, 12(5), 1361–1395.

Gelman, A., and G. King (1990): "Estimating incumbency advantage without bias," *American Journal of Political Science*, pp. 1142–1164.

Gerardi, D., and L. Yariv (2007): "Deliberative voting," *Journal of Economic Theory*, 134(1), 317–338.

Gershkov, A., J. K. Goeree, A. Kushnir, B. Moldovanu, and X. Shi (2013): "On the equivalence of Bayesian and dominant strategy implementation," *Econometrica*, 81(1), 197–220.

Hummel, P. (2012): "Deliberation in large juries with diverse preferences," *Public Choice*, 150(3-4), 595–608.

Imbens, G., and K. Kalyanaraman (2012): "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *Review of Economic Studies*, 79(3), 933–959.

Imbens, G. W., and T. Lemieux (2008): "Regression discontinuity designs: A guide to practice," *Journal of Econometrics*, 142(2), 615–635.

Ishida, J., and T. Shimizu (2016): "Cheap talk with an informed receiver," *Economic Theory Bulletin*, 4(1), 61–72.

Kellerer, H., U. Pferschy, and D. Pisinger (2004): *Knapsack problems.* Springer Science & Business Media.

Ledyard, J., D. Porter, and R. Wessen (2000): "A market-based mechanism for allocating space shuttle secondary payload priority," *Experimental Economics*, 2(3), 173–195.

Lee, D. S. (2008): "Randomized experiments from non-random selection in US House elections," *Journal of Econometrics*, 142(2), 675–697.

Lee, D. S., and T. Lemieux (2010): "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 48, 281–355.

Levitt, S. D., and C. D. Wolfram (1997): "Decomposing the sources of incumbency advantage in the US House," *Legislative Studies Quarterly*, pp. 45–60.

Li, S. (2015): "Obviously Strategy-Proof Mechanisms," *Available at SSRN 2560028.*

Liang, C.-Y. (2013): "Is there an incumbency advantage or cost of ruling in proportional election systems?," *Public Choice*, 154(3-4), 259–284.

Lizzeri, A., and N. Persico (2005): "A drawback of electoral competition," *Journal of the European Economic Association*, pp. 1318–1348.

Luton, R., and R. P. McAfee (1986): "Sequential procurement auctions," *Journal of Public Economics*, 31(2), 181–195.

Manelli, A. M., and D. R. Vincent (2010): "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model," *Econometrica*, 78(6), 1905–1938.

Maskin, E. (2002): "How to Reduce Greenhouse Gas Emissions: An Application of Auction Theory," *Nancy L. Schwartz Memorial Lecture, Northwestern University.*

McCrary, J. (2008): "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, 142(2), 698–714.

Milesi-Ferretti, G. M., R. Perotti, and M. Rostagno (2002): "Electoral systems and public spending," *Quarterly Journal of Economics*, pp. 609–657.

Milgrom, P., and I. Segal (2014): "Deferred-acceptance auctions and radio spectrum reallocation," in *Proceedings of the 15th ACM Conference on Economics and Computation.*

Mookherjee, D., and S. Reichelstein (1992): "Dominant strategy implementation of Bayesian incentive compatible allocation rules," *Journal of Economic Theory*, 56(2), 378–399.

Moulin, H. (1999): "Incremental cost sharing: Characterization by coalition strategy-proofness," *Social Choice and Welfare*, 16(2), 279–320.

Moulin, H., and S. Shenker (2001): "Strategyproof sharing of submodular costs:budget balance versus efficiency," *Economic Theory*, 18(3), 511–533.

Mu'Alem, A., and N. Nisan (2008): "Truthful approximation mechanisms for restricted combinatorial auctions," *Games and Economic Behavior*, 64(2), 612–631.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, pp. 58–73.

Pai, M. M., and R. Vohra (2014): "Optimal auctions with financially constrained buyers," *Journal of Economic Theory*, 150, 383–425.

Persson, T., G. Roland, and G. Tabellini (2007): "Electoral Rules and Government Spending in Parliamentary Democracies," *Quarterly Journal of Political Science*, 2, 155–188.

Persson, T., and G. Tabellini (2004): "Constitutions and economic policy," *Journal of Economic Perspectives*, pp. 75–98.

Pettersson-Lidbom, P. (2008): "Do Parties Matter for Economic Outcomes? A Regression-Discontinuity Approach," *Journal of the European Economic Association*, 6(5), 1037–1056.

——— (2012): "Does the size of the legislature affect the size of government? Evidence from two natural experiments," *Journal of Public Economics*, 96(3), 269–278.

Singer, Y. (2010): "Budget feasible mechanisms," in *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pp. 765–774. IEEE.

Snyder, J. M., O. Folke, and S. Hirano (2015): "Partisan imbalance in regression discontinuity studies based on electoral thresholds," *Political Science Research and Methods*, 3(02), 169–186.

Strausz, R. (2003): "Deterministic mechanisms and the revelation principle," *Economics Letters*, 79(3), 333–337.

Thordal-Le Quement, M., and V. Yokeeswaran (2015): "Subgroup deliberation and voting," *Social Choice and Welfare*, 45(1), 155–186.

Trounstine, J. (2011): "Evidence of a local incumbency advantage," *Legislative Studies Quarterly*, 36(2), 255–280.

Uppal, Y. (2010): "Estimating Incumbency Effects in the US State Legislatures: A Quasi-Experimental Study," *Economics & Politics*, 22(2), 180–199.

US Department of the Treasury (2011): "Treasury Announces Intent to Sell Warrant Positions in Public Dutch Auctions," *Press Release*, `http://www.treasury.gov/press-center/press-releases/Pages/tg1193.aspx`, accessed: February 2, 2015.

Curriculum Vitae

2010-2016	PhD Studies in Economics, Center for Doctoral Studies in Economics (CDSE), Graduate School of Economic and Social Sciences (GESS), University of Mannheim. Thesis: *Essays in Collective Decision Making*
2012	M.Sc. Economics, University of Mannheim.
2011-2012	Visiting Scholar, Yale University.
2007-2010	B.Sc. Economics, University of Mannheim.
2006-2007	Community service in lieu of military service (Zivildienst).
2006	High School Diploma (Abitur), Gymnasium Andreanum, Hildesheim.