Five Essays in Economic Theory

Inaugural-Dissertation

zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften

durch die

Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von

Moritz Drexl

aus Hamburg

Bonn 2014

Erstreferent: Prof. Dr. Benny Moldovanu

Zweitreferent: Prof. David C. Parkes, PhD

Tag der mündlichen Prüfung: 24. September 2014

Acknowledgements

I am grateful to many people for their support in preparing this thesis.

First, I wish to thank my supervisor Benny Moldovanu for providing insightful hints and comments. I also want to thank David Parkes for his guidance during my stay abroad and for all the interesting discussions we had.

Special thanks to my co-author and friend Andy Kleiner with whom it is a real pleasure to share an office, do research and write papers.

Further, I want to thank the numerous people I've talked to and whose comments helped to improve this thesis, including Gabriel Carroll, Drew Fudenberg, Jerry Green, Martin Hellwig, Werner Hildenbrand, Daniel Krähmer, Eric Maskin, David Miller, Michael Ruberry, Ilya Segal, Alexander Teytelboym, Alexander Westkamp, Zaifu Yang, as well as all the participants in the seminars and talks where parts of this thesis were presented.

The Bonn Graduate School of Economics provided financial support, for which I am very grateful. In particular, I thank Silke Kinzig, Pamela Mertens and Urs Schweizer for their endless efforts in providing an excellent research environment.

Finally, I wish to thank my friends and my family, and especially my girlfriend Janina for her love and patience.

Contents

Introduction

This thesis covers two main areas of microeconomic theory. The first three chapters present the results of joint research with Andreas Kleiner, and are contributions to the theory of mechanism design. The last two chapters contribute to the literature on general equilibrium in markets with indivisible goods.

First Part, Chapters One to Three

Mechanism design is concerned with the implementability of social choice functions when agents are privately informed about their preferences. A social choice function is a rule which specifies how to choose among a given set of alternatives, for each possible combination of preferences that a population may have over these alternatives. Since preferences are private information, it is reasonable that the agents will not necessarily reveal their true preference when asked for it, in order to apply the social choice function. However, in many situations a mechanism can be designed to solve this problem. A mechanism specifies the rules of a game such that, if the agents with certain preferences play an equilibrium of the game, the outcome is precisely that which is prescribed by the social choice function for these preferences. Then the mechanism is said to be incentive compatible and to implement the social choice function.

The theory of mechanism design aims at characterizing the set of social choice functions that are implementable with respect to certain notions of equilibrium, and then optimizing over this set of functions according to different objective functions and subject to additional constraints. For example, an auction can be interpreted as a social choice function and one can ask for the auction that yields the highest revenue, provided that the participants bid in a Bayes-Nash equilibrium. The answer is that the seller should conduct a second-price auction with a reserve price that depends on how the seller estimates the bidders' preferences to be distributed (Myerson 1981).

The first three chapters of this thesis focus on a particular objective function for the determination of an optimal mechanism, and study it in three different settings. In each setting, we identify mechanisms that maximize expected residual surplus. This is the aggregate utility (or welfare) of all agents and therefore explicitly includes monetary transfers that are possibly needed in order to make the mechanism incentive compatible. This contrasts most of the literature on mechanism design which does not consider as welfare-reducing the transfers that leave the group of agents (sometimes also referred to as money burning). Underlying the computation of expected residual surplus is an assumed distribution of preferences which we always require to satisfy monotone hazard rates. This is a widely used assumption in the mechanism design literature. In all three chapters we will further require implementation in ex-post equilibrium which means that the agents' strategies remain optimal even when they know the preferences of the other players. This ensures that the mechanisms are robust to informational disturbances and is helpful in guiding practical decisions on which social choice rule to pick. Using this approach, we can also explain the prevalence of certain mechanisms in practice.

Specifically, in the first chapter we look at settings in which a group of agents is faced with the decision to accept or reject a given proposal. This can be, for instance, the decision to pass or reject a bill, or whether to hire a new colleague. Every member of the group has a (privately known) positive or negative willingness-to-pay for the proposal. While the efficient decision rule would be to accept the proposal if and only if the average willingness-to-pay is positive, this can only be implemented if transfers leave the group of agents. Our mechanism design approach of identifying the social choice rule that maximizes residual surplus establishes that the best mechanism for this setting is a simple majority voting rule which does not involve transfers at all. This is in line with the fact that in most practical situations the decision is carried out without the use of transfers and therefore we provide a rationale for the widespread use of voting.

The second chapter considers a dynamic version of the above setting. In every period, the group of agents has to decide whether to accept or reject a different proposal. Although we assume that utility is not transferable (i.e., money is not feasible, usually for ethical or other reasons), a dynamic social choice rule may condition on past decisions and behavior. This enables the modeling of phenomena like vote trading or explicit mechanisms like budgeted veto rights. The main insight of this chapter is that changes of the mechanism in future periods that depend on present behavior affect an agent's incentives in the same way as monetary transfers, which are usually used to align incentives. For example, if an agent exercises his veto right today, he will not have it in future periods, which changes his expected future utility. This interpretation of expected future utility as monetary transfers allows us to apply similar techniques as in the first chapter, and we can derive the main result that the welfare-optimal dynamic decision rule in every period decides according to the same majority voting rule. This implies that the outcome of vote trading games or veto rights mechanisms is welfare-inferior to periodic majority voting.

The third chapter studies the allocation of a private good among two agents in the context of residual surplus maximization. This is done in two different settings: In the auction setting, the good does not initially belong to any of the agents. We derive that any optimal mechanism takes one of two simple forms. Either it is a posted price mechanism, where the good is given to one of the agents unless both agents agree to trade the good at a prespecified strike price. Or it is an option mechanism, where the good is given to one of the agents and the other agent is given the option to buy the good from the first agent for a prespecified strike price. The second setting is bilateral trade, where one agent (the seller) initially owns the good. Here, we can show that posted price mechanisms are optimal trading mechanisms. Since the optimal mechanism has a balanced budget, this result shows that in the traditional literature on bilateral trade, budget-balancedness does not need to be imposed a priori (Myerson and Satterthwaite 1983, Hagerty and Rogerson 1987).

Second Part, Chapters Four and Five

The second part of this thesis studies the existence and computation of market equilibria in exchange economies with indivisible goods. In these models, agents from a given population have quasi-linear preferences over bundles of certain goods as well as money. A competitive equilibrium (sometimes also called a market equilibrium or Walrasian equilibrium) consists of a price for each good, such that the market clears when every agent demands its most-preferred bundle at these prices. Next to the question of the existence of such a market equilibrium, a central concept for the study of exchange economies is a tâtonnement process that adjusts prices until an equilibrium is attained. Underlying such a process is the idea of a Walrasian auctioneer that changes ask prices in response to supply and demand (Walras 1874), and indeed tâtonnement processes are closely linked to iterative auction formats.

In the context of indivisible goods, a central assumption on the set of possible preferences is that of gross substitutes. An agent having gross substitutes preferences views all goods as substitutes for each other, in the sense that, if the price for one of the goods increases, he will buy (weakly) more of every other good. This assumption ensures that competitive equilibria always exist in the standard auction environment (Kelso and Crawford 1982), where a set of items is available for sale to a group of potential buyers. This environment, as well as the set of possible preferences, have since then been generalized. For example, the two-sided structure considered in the auction setting can be extended to a network of trading relationships (Hatfield, Kominers, Nichifor, Ostrovsky and Westkamp 2013).

The fourth chapter of this thesis provides a generalization of the gross substitutes condition in this trading network environment. In these economies, agents are located at nodes in a network and can engage in various trading relationships with their "neighbors" in which they are either the seller or the buyer. Underlying the adjacent trades are goods over which agents have quasi-linear preferences. We assume that the preferences satisfy the following assumption (gross substitutes and complements, see Sun and Yang 2006): The set of possible trades can be divided into two sets which can be thought of as tables and chairs. Agents view goods in one set as substitutes for each other (so one table substitutes another table), but view goods in different sets as complementing each other (so every table complements every chair). We show that with this class of preferences, a competitive equilibrium always exists and thereby unify previous generalizations of the Kelso and Crawford model.

By harnessing the interpretation of the gross substitutes condition as a geometric property of discrete convex functions, the final chapter of this thesis studies the connections between tâtonnement processes for economies with indivisible goods and algorithms for the minimization of discrete convex functions. Specifically, for a valuation function that represents the preferences of an agent, the indirect utility function can be considered. This is the utility an agent gets if he chooses his most-preferred bundle, given prices for each good, and mathematically corresponds to the convex conjugate of the agent's valuation function. If the agent has gross substitutes preferences, then the valuation function as well as the indirect utility function belong to classes of discrete convex functions with nice combinatorial structure. For these functions, market equilibria correspond exactly to the set of prices that minimize aggregate indirect utility, and steepest descent algorithms can be used to find these prices. Using these connections, we are able to generalize existing tâtonnement processes (Ausubel 2006) to arbitrary exchange economies with agents that are buyers and/or sellers of multiple units of different goods. These results are applied to obtain price adjustment processes for the trading network economies treated in the fourth chapter, as well as for models with gross substitutes and complements preferences.

Chapter 1

Why Voting? A Welfare Analysis

Voting is commonly applied in collective decision making, but at the same time it is criticized for being inefficient. We address this apparent conflict and consider committees deciding collectively between accepting a given proposal and maintaining the status quo. Committee members are privately informed about their valuations and monetary transfers are possible. We solve for the social choice function maximizing utilitarian welfare, which takes monetary transfers to an external agency explicitly into account. For regular distributions of preferences, we find that it is optimal to exclude monetary transfers and to decide by qualified majority voting.

1. Introduction

Why is voting predominant in collective decision making? A common view is that often it is immoral to use money. This view is plausible, for example, when deciding who should receive a donated organ or whether a defendant should be convicted. However, it explains less convincingly why shareholders vote on new directors at the annual meeting, why managing boards of many companies make important operative decisions by voting, or why hiring committees vote when deciding on a new appointment. Indeed, voting is criticized for its inefficiency, and the economic literature argues that collective decisions can be improved if transfers are used to elicit preference intensities. But redistributing these transfers within the group introduces incentive problems, while wasting them reduces welfare. We model these considerations explicitly, and show that voting maximizes welfare.

Our analysis closely follows standard models of collective decision making: A finite population of voters decides collectively whether to accept a given proposal or to maintain the status quo. Agents are privately informed about their valuations and have quasi-linear utilities. Monetary transfers are feasible as long as they create no budget deficit and agents are willing to participate in the decision process. In contrast to much of the literature, we consider a utilitarian welfare function that takes monetary transfers to an external agency into account. We then investigate which strategy-proof social choice function maximizes this aggregate expected utility.

Our main result is that the optimal anonymous social choice function is implementable by qualified majority voting. Under such schemes, agents simply indicate whether they are in favor or against the proposal, and the proposal is accepted if the number of agents being in favor is above a predetermined threshold. This implies that, even though it is possible to use monetary transfers, it is optimal *not* to use them. Specifically, we show that any anonymous decision rule that relies on monetary transfers wastes money to such an extent that it is inferior to voting. It follows that it is *not* possible to improve upon voting without giving up reasonable properties of the social choice function. Our result thereby justifies the widespread use of voting rules in practice, and provides a link between mechanism design theory and the literature on political economy.

Our finding that voting performs well from a welfare perspective stands in sharp contrast to the previous literature, which suggests to implement the value-maximizing public decision. However, this does not achieve the first-best because it induces budget imbalances (see, e.g., Green and Laffont 1979). While it is traditionally assumed that money wasting has no welfare effects, we consider a social planner that cares about aggregate transfers. This approach seems reasonable for at least two reasons: First, a social planner might be interested in implementing the decision rule that maximizes the agents' expected utility, which in turn depends on the payments they have to make. Second, groups often choose the rule by which they decide themselves, and when making this choice they take the payments they have to make into account. Hence, our approach provides an explanation for which decision rules are likely to prevail in practice.

Our result, that transfer-free voting schemes dominate more complex decision rules, follows from two basic observations. In a first step, we analyze the transfers that are necessary to implement a given decision rule. Incentive compatibility fixes the payment function up to a term that only depends on the reports of all other agents. We show that the requirements of (a) no money being injected and (b) all agents being willing to participate in the decision procedure, entirely fix the payment functions for any anonymous decision rule. In particular, it turns out that if money is necessary to induce truthful reporting then it has to be wasted. As an application, this implies that any anonymous social choice function is implementable with a balanced budget if and only if it can be implemented by qualified majority voting. In a second step, we then analyze the trade-off between increasing efficiency of the public decision and reducing the waste of monetary resources. For regular distribution functions, we show that this trade-off is solved optimally by not using money at all. This implies that the optimal social choice function is implementable by qualified majority voting. We also characterize the minimum number of votes that is optimally required for the adoption of the proposal.

Related Literature

Formal analyses of the question "should we use monetary transfers or not?" are rare; to the best of our knowledge, the only attempts are arguments that voting mechanisms are easy and perform well for large populations (Ledyard and Palfrey 2002), and that voting rules are coalition-proof (Bierbrauer and Hellwig 2012). We complement these papers by arguing that voting is optimal from a utilitarian perspective.

The fact that the optimal decision scheme does not use transfers relates our work to the analysis of optimal collective decision rules when monetary transfers are not feasible. This literature was initiated by Rae (1969), who compares utilitarian welfare of different voting rules and shows that simple majority voting (where a proposal is accepted if at least half of the population votes for it) is optimal if preferences are symmetric across outcomes. Recently, this approach was generalized to include more general decision rules (Azrieli and Kim 2012), to allow for correlated valuations (Schmitz and Tröger 2012) and to consider more than two alternatives (Gershkov, Moldovanu and Shi 2013).

Barbera and Jackson (2004) study a model where agents not only vote on a given proposal, but in a first stage decide on which voting rule to use in the second stage. They argue that only "self-stable" rules, i.e., voting rules that would not be changed once in place, are likely to prevail. If agents are ex-ante symmetric, only voting rules that maximize utilitarian welfare satisfy this condition. We contribute to this branch of the literature by showing that, in our setting, the exclusion of money is not costly.

Our insight that monetary transfers are not necessarily welfare-increasing relates our work to studies that exclude monetary transfers but allow for costly signaling. These studies assume that signaling efforts are wasteful and cannot be redistributed. It is shown that the welfare-maximizing allocation of private goods relies only on prior information and completely precludes wasteful signaling (Hartline and Roughgarden (2008), Yoon (2011), Condorelli (2012), Chakravarty and Kaplan (2013); see McAfee and McMillan (1992) for a result in a similar vein). In contrast, we allow for monetary transfers from and between agents and show that in a public good setting similar economic trade-offs arise.

An extensive literature in mechanism design studies allocation problems when monetary transfers are feasible. While VCG mechanisms implement the value-maximizing public decision (Groves 1973), this comes at the cost of budget imbalances that cannot be redistributed without distorting incentives (Green and Laffont 1979, Walker 1980).[1] Therefore, these mechanisms achieve the first-best only under the assumption that the social planner does not care about monetary resources. An opposite approach, where the budget is required to be exactly balanced, is pursued in Laffont and Maskin (1982).

The budget imbalances of VCG mechanisms might be less severe if they were quantitatively negligible in practical applications. This argument has been put forward by Tideman and Tullock (1976), who conjecture that wasted transfers are not important for large populations[2] and VCG mechanisms therefore approximate the first-best. In Section 4 we discuss how our result relates to this observation.

A small part of the literature, which also considers money burning to be welfare-reducing, studies the allocation of a private good. Miller (2012) shows that the optimal mechanism never allocates efficiently and in some cases wastes monetary resources. If there are only two agents and the distribution functions are regular then the optimal mechanism transfers money and has a balanced budget (Drexl and Kleiner 2012, Shao

[1]For an approach using a weaker equilibrium concept see d'Aspremont and Gerard-Varet (1979). Note that the equivalence between dominant strategy and Bayes-Nash incentive compatible mechanisms established by Gershkov, Goeree, Kushnir, Moldovanu and Shi (2013) does not hold in this model as the budget is constrained ex-post.

[2]This claim was formally verified by Green and Laffont (1977).

and Zhou 2012). In contrast, the optimal social choice function in the present chapter does not use money.

Finding the optimal social choice function involves understanding which part of the payments can be redistributed without distorting incentives (see also the work of Cavallo 2006). Our focus on anonymous social choice functions for a public good setting allows us to solve this problem.

We proceed as follows: We present the model in Section 2, derive our main result in Section 3 and provide a short discussion of the result in Section 4.

2. Model

We consider a population of N agents[3] deciding collectively on a binary outcome $X \in \{0,1\}$. We interpret this as agents deciding whether they accept a proposal (in which case $X = 1$) or reject it and maintain the status quo ($X = 0$). Given a collective decision X, the utility of agent i is given by $\theta_i \cdot X + t_i$, where θ_i is the agent's valuation for the proposal and t_i is a transfer to agent i.[4] Each agent is privately informed about his valuation, which is drawn independently from a type space $\Theta := \left[\underline{\theta}, \overline{\theta}\right]$ according to a distribution function F with positive density f. To make the problem interesting we assume that $\underline{\theta} < 0 < \overline{\theta}$.[5] Both type space and distribution function are common knowledge. Let Θ^N denote the product type space consisting of complete type profiles with typical element $\theta = (\theta_i, \theta_{-i})$.

A social choice function in this setting determines for which preference profiles the proposal is accepted and which transfers are made to the agents. Formally, a *social choice function* is a pair $G = (X^G, T^G)$ consisting of a *decision rule*

$$X^G \; : \; \Theta^N \; \rightarrow \; \{0,1\}$$

and a *transfer rule*

$$T^G : \Theta^N \rightarrow \mathbb{R}^N$$

such that, for any realized preference profile θ, $X^G(\theta)$ is the decision on the public outcome and $T_i^G(\theta)$ is the transfer received by agent i. A social choice function is *feasible* if, for any realization of preferences, no injection of money from an external agency is necessary, i.e., if

$$\sum_{i \in N} T_i^G(\theta) \leq 0. \tag{F}$$

In many situations agents have the outside option to abstain from the decision process and leave the decision to the other agents. It is then without loss of generality

[3]For convenience, we also write N for the set of agents $\{1, \ldots, N\}$.

[4]Our analysis applies to costless projects as well as to costly projects with a given payment plan, in which case the valuation of agent i is interpreted as her net valuation taking her contribution into account. Also note that the analysis accommodates more general utility functions: Take any quasi-linear utility function such that the utility difference between $X = 1$ and $X = 0$ is continuous and strictly increasing in θ_i. Redefining the type to equal the utility difference, we can proceed with our analysis without change.

[5]The analysis directly extends to cases where $\underline{\theta} = -\infty$ and/or $\overline{\theta} = \infty$.

to consider social choice functions that ensure participation in the following sense: If agent i leaves the decision process, the social choice function chooses some alternative $\underline{X}_i(\theta_{-i})$. Then the social choice function satisfies *universal participation* (see, e.g., Green and Laffont 1979) if, given this outside option, all agents prefer to participate in the decision process:[6]

$$\theta_i X^G(\theta) + T_i^G(\theta) \geq \theta_i \underline{X}_i^G(\theta_{-i}). \tag{UP}$$

This constraint is weaker than the requirement that every agent derive utility of at least zero (often called individual rationality). For instance, majority voting satisfies universal participation but in general it is not individually rational.

Definition 1. *We call a decision rule X^G* anonymous *if it is independent of the agents' identities, i.e. if, for each permutation $\pi : N \to N$ and corresponding function $\hat{\pi}(\theta) = (\theta_{\pi(1)}, \ldots, \theta_{\pi(N)})$, it holds that $X^G(\theta) = X^G(\hat{\pi}(\theta))$ for all θ.*

A social choice function is anonymous if the associated decision rule is anonymous.

This is a weak notion of anonymity, requiring only that the names of the agents do not affect the public decision. However, focusing on anonymous social choice functions is a potentially severe restriction.[7] Nonetheless, it is often reasonable to impose anonymity as many fairness concepts build on this assumption (e.g., equal treatment of equals). This requirement also has a long tradition in social choice theory, see for example, Moulin (1983).[8]

We are interested in social choice functions that are *strategy-proof*, i.e., for which there exists a mechanism and an equilibrium in dominant strategies for the strategic game induced by this mechanism such that, for any realized type profile, the equilibrium outcome corresponds to the outcome that the social choice function stipulates. Requiring social choice functions to be strategy-proof is a standard approach in social choice theory (see, e.g., Moulin 1983).[9]

Throughout the chapter we focus on anonymous and feasible social choice functions that are strategy-proof and satisfy universal participation. Which social choice function should a utilitarian planner choose? Given that the value-maximizing decision cannot be implemented with a balanced budget, a utilitarian planner should implement the

[6]We note that our analysis does not depend on any particular form of the function $\underline{X}_i$. This outside option could also depend on the privately observed valuation of agent i without any change in the analysis.

[7]For example, it excludes the use of "sampling Groves mechanisms" (Green and Laffont 1979), where a VCG mechanism is used for a subset of the population and the budget surplus is redistributed to non-sampled agents.

[8]Note that this assumption would be without loss of generality if we allowed for stochastic decision rules. Given any social choice function (X^G, T^G), apply this function after randomly permuting the agents. This defines a new social choice function $(\tilde{X}^G, \tilde{T}^G)$ that is anonymous and achieves the same utilitarian welfare. While this new rule treats all agents equally ex-ante, it is possible that agents with the same valuations are treated very differently after the uncertainty about the randomization is resolved.

[9]Bierbrauer and Hellwig (2012) show for the model we consider that strategy-proofness is equivalent to robust implementation in the spirit of Bergemann and Morris (2005).

second-best, i.e., maximize utilitarian welfare given by

$$U\left(X^G, T^G\right) := \mathbb{E}_\theta \left[\sum_{i=1}^{N} \left[\theta_i X^G(\theta) + T_i^G(\theta) \right] \right],$$

where the expectation is taken with respect to the prior distribution of θ. The assumption that the planner perfectly knows the prior distribution of types, although being very common in the literature on mechanism design, might be too strong in some settings. Note however, that the optimal social choice function derived in Theorem 1 does not depend on the exact distribution of types. Moreover, as we focus on robust implementation, misspecifications do not affect incentives and hence the performance of the optimal social choice function is not very sensitive to slight misestimations of the distribution of types.

3. Results

To implement a given social choice function, we invoke the revelation principle (Gibbard 1973). It follows that we can focus without loss of generality on direct revelation mechanisms in which it is a dominant strategy for agents to report their valuations truthfully. Hence, a mechanism is given by a tuple (x, t), where $x : \Theta^N \to \{0, 1\}$ maps reported types into a collective decision and, for each agent i, $t_i : \Theta^N \to \mathbb{R}$ maps reported types into the payment received by that agent. The requirement that a social choice function be strategy-proof translates to

$$\theta_i x(\theta_i, \theta_{-i}) + t_i(\theta_i, \theta_{-i}) \geq \theta_i x(\hat{\theta}_i, \theta_{-i}) + t_i(\hat{\theta}_i, \theta_{-i}) \quad \text{for all} \quad \theta_{-i}, \theta_i, \hat{\theta}_i. \tag{IC}$$

A mechanism is *qualified majority voting (with threshold k)*, if $x(\theta) = 1$ if and only if $|\{i : \theta_i \geq 0\}| \geq k$ and if in no case monetary transfers are made, i.e., $t_i(\theta) = 0$ for all i and θ.

Definition 2. *A distribution function F has* monotone hazard rates *if the hazard rate $\frac{f(\theta_i)}{1-F(\theta_i)}$ is non-decreasing in θ_i for $\theta_i \geq 0$ and the reversed hazard rate $\frac{f(\theta_i)}{F(\theta_i)}$ is non-increasing in θ_i for $\theta_i \leq 0$.*

This assumption is well-known from the literature on optimal auctions and procurement auction design; it is satisfied by many commonly employed distribution functions, for example by the uniform, (truncated) normal, and exponential distributions.

We are now ready to state our main result.

Theorem 1. *Suppose F has monotone hazard rates. Then the optimal social choice function is implementable by qualified majority voting with threshold $\lceil k \rceil$, where*

$$k := \frac{-N \, \mathbb{E}[\theta_i | \, \theta_i \leq 0]}{\mathbb{E}[\theta_i | \, \theta_i \geq 0] - \mathbb{E}[\theta_i | \, \theta_i \leq 0]}.$$

That is, the optimal decision rule does not rely on monetary transfers at all and can be implemented using a simple indirect mechanism where each agent indicates whether she is in favor of or against the proposal. It is accepted if more than $\lceil k \rceil$ voters are in favor.[10] The following example illustrates how voting mechanisms compare to the first-best and the best VCG mechanism.

Example 1. *Let $N = 2$ and θ_i be independently and uniformly distributed on $[-3, 3]$ for $i = 1, 2$. If valuations were publicly observable the first-best could be implemented, which would yield welfare $U_{FB} = \frac{1}{2}\mathbb{E}[\theta_1 + \theta_2 \mid \theta_1 + \theta_2 \geq 0] = 1$. The best VCG mechanism is the pivotal mechanism, which gives welfare $U_{VCG} = \frac{1}{2}$ (see the Appendix). In contrast, unanimity voting, that is, accepting the proposal if and only if both agents have a positive valuation, yields welfare $U_{UV} = \frac{1}{4}\mathbb{E}[\theta_1 + \theta_2 \mid \theta_1 \geq 0, \theta_2 \geq 0] = \frac{3}{4}$. Hence, the welfare loss due to private information is twice as large under the best VCG mechanism as compared to unanimity voting.*

The broader implications of Theorem 1 are discussed in Section 4 and a formal proof is provided in the Appendix. In the following, we build some intuition for this result.

As a first step, Lemma 1 characterizes direct mechanisms that are strategy-proof. It shows that the transfer of every type is determined by the decision rule up to a term that only depends on the reports of the other agents. Since this term changes the transfers of an agent without affecting his incentives, we call it "redistribution payment."

As a second step, we show that, for any anonymous social choice function, positive redistribution payments are not feasible and therefore all collected payments have to be wasted (Lemma 2). In general, it is easy to build strategy-proof and budget-balanced social choice functions by ignoring one agent in the public decision and awarding him all payments by the other agents. Anonymity not only rules out this possibility, but one can prove that any mechanism which has positive redistribution payments is necessarily asymmetric.

Given that money cannot be redistributed in anonymous social choice functions, there is a direct trade-off between improving the decision rule and reducing the outflow of money. We show, as a third step, that this conflict is resolved optimally in favor of no money burning. To gain some intuition, fix a type profile of the other agents, θ_{-i}. Strategy-proofness implies that there is a cutoff θ_i^* such that the proposal will be accepted if the type of agent i is above θ_i^*. To solve for the optimal decision rule we need to find the optimal cutoff. Assume that the sum of valuations $\sum_{j \neq i} \theta_j + \theta_i^*$ is negative. Marginally increasing the cutoff leads to a rejection of the proposal which in this case increases efficiency (with a positive effect on welfare proportional to $f(\theta_i^*)$). On the other hand, strategy-proofness implies that agents with a type above the cutoff make a payment equal to the cutoff. Increasing the cutoff increases these payments (with a corresponding negative effect on welfare proportional to $1 - F(\theta_i^*)$). Monotone hazard rates imply that if the positive effect outweighs the negative effect at θ_i^* and therefore it is beneficial to marginally increase the cutoff, then it is optimal to set the cutoff to the highest possible value. Symmetric arguments imply that it is optimal to set all cutoffs either equal to zero or to the boundary of the type space, and hence that the optimal mechanism can be implemented by a voting rule.

Finally, the optimal number of votes required in favor of a proposal is given by the smallest integer number k such that the expected aggregate welfare of a proposal,

[10] See also Nehring (2004), Barbera and Jackson (2006).

given that k out of N voters have a positive valuation, is positive. Hence, the optimal threshold required for qualified majority voting depends on the conditional expected values given that the valuation is either positive or negative. Simple majority voting is optimal if valuations are distributed symmetrically around 0. If, however, opponents of a proposal are expected to have a stronger preference intensity, then it is optimal to require a qualified majority that is larger than simple majority.

As an easy consequence, Lemma 1 and Lemma 2 permit a characterization of the set of strategy-proof social choice functions that have a balanced budget.

Corollary 1. *A feasible and anonymous social choice function satisfying universal participation has a balanced budget if and only if it is implementable by qualified majority voting.*

In comparison to this corollary, Theorem 1 allowed for a larger class of social choice functions that potentially waste money. While we determine the optimal social choice function in this larger class in the theorem, this corollary characterizes any implementable social choice function in the smaller class of budget-balanced social choice functions. A closely related result has been obtained by Laffont and Maskin (1982), who in addition require weak Pareto efficiency but do not impose participation constraints.

4. Discussion

This chapter shows that utilitarian welfare, which takes transfers into account, is maximized by using qualified majority voting. Our result resolves the apparent conflict between the widespread use of such mechanisms in practice and the intuition that accounting for preference intensities can improve collective decisions. In particular, we show that the costs of accounting for preference intensities outweigh the benefits and the VCG mechanism is inferior to voting. In contrast, Tideman and Tullock (1976) argue that payments vanish as the number of agents gets large and hence the VCG mechanism should be used instead of voting. However, while it is generically true that the VCG mechanism approximates the first-best if the population is large enough, this is not sufficient for being superior to voting. In fact, voting also approximates the first-best. Moreover, for any fixed population, it turns out that voting provides a higher expected welfare. More generally, Theorem 1 indicates that being welfare-inferior to voting is not a problem of the VCG mechanism, but that it is in fact not possible to improve upon voting under the normative requirements of robust implementation and equal treatment of equals.

Classical social choice theory suggests that decisions should depend on the average willingness-to-pay in the population, i.e., a proposal should be accepted if the average willingness-to-pay is positive. In contrast, decision rules considered in political economy and implemented in practice typically depend only on the number of agents with a positive willingness-to-pay. By taking an optimal mechanism design approach we are able to reconcile mechanism design theory with social choice practice and the literature on political economy.

An important question in this respect concerns the robustness of our results to alternative specifications of the decision problem. First, if one considers more general problems with more than two possible outcomes, the results will crucially depend on the restrictions imposed on preferences.[11] Second, it would be interesting to relax some of the restrictions we imposed on the social choice functions. While it appears that relaxing universal participation does not change the spirit of our results, our analysis depends crucially on the assumption of anonymity.

Appendix

Verification of Example 1. Welfare of the pivot mechanism can be expressed as the difference between the welfare of the first-best and the transfers needed to implement the efficient decision:

$$U_{VCG} = U_{FB} - \frac{4}{36}\int_{-3}^{0}\int_{0}^{-\theta_1} (-\theta_2)\; d\theta_2\; d\theta_1 = \frac{1}{2}$$

Here, we used the fact that transfers are symmetric in the four regions $\{\theta \mid \theta_i \geq 0, \theta_j \leq 0, \theta_i + \theta_j \lesseqgtr 0\}$ and zero everywhere else. $\square$

The following lemma is a standard characterization of strategy-proof mechanisms.

Lemma 1. *A mechanism is strategy-proof if and only if, for each agent i,*

1. *$x(\theta_i, \theta_{-i})$ is non-decreasing in θ_i for all θ_{-i} and*
2. *there exists a function $h_i(\theta_{-i})$, such that for all θ,*

$$\theta_i x(\theta_i, \theta_{-i}) + t_i(\theta_i, \theta_{-i}) = h_i(\theta_{-i}) + \int_0^{\theta_i} x(\beta, \theta_{-i}) d\beta. \tag{1}$$

Equation (1) suggest the following definition:

Definition 3. *Agent i is* pivotal *at profile θ, if $\theta_i x(\theta) \neq \int_0^{\theta_i} x(\beta, \theta_{-i}) d\beta$.*

A necessary condition for agent i to be pivotal at θ is that $x(\theta) \neq x(0, \theta_{-i})$. If agent i is not pivotal at a given profile (θ_i, θ_{-i}) then her payment equals $h_i(\theta_{-i})$. If she is pivotal at this profile, her transfer is reduced by $\theta_i x(\theta) - \int_0^{\theta_i} x(\beta, \theta_{-i}) d\beta$.

Lemma 2. *Suppose a mechanism (x, t) is anonymous. Then $h_i(\theta_{-i}) = 0$ for all i and θ_{-i}.*

Proof. The proof consists of two steps.

Step 1: For all i and θ_{-i}, there exists θ_i such that no agent is pivotal at (θ_i, θ_{-i}).

Note that all agents that are pivotal at profile θ submit reports of the same sign: If $x(\theta) = 1$ then monotonicity implies that $x(0, \theta_{-i}) = 1$ for all agents i with $\theta_i < 0$ and hence only agents with positive reports can be pivotal (and similarly for $x(\theta) = 0$).

[11] For example, for quadratic utilities and a continuum of alternatives, the efficient allocation rule can be implemented with a balanced budget (Groves and Loeb 1975).

Fix an arbitrary agent i and a report profile $\theta_{-i} \in \Theta^{N-1}$. Suppose without loss of generality that $x(0, \theta_{-i}) = 1$ and that all agents that are pivotal at $(0, \theta_{-i})$ submit positive reports (if no agent is pivotal at this profile, we are done; if $x(0, \theta_{-i}) = 0$ analogous arguments hold). We show that no agent is pivotal at profile $\theta := (\theta_{j^*}, \theta_{-i})$, where $j^* \in \arg\max_j \theta_j$. Monotonicity implies that $x(\theta) = x(0, \theta_{-i}) = 1$ and hence agent i is not pivotal. Anonymity implies that agent j^* is not pivotal. The claim is proved if we can show that if j is not pivotal at θ and $\theta_{j'} \leq \theta_j$, then j' is not pivotal at θ. Assume to the contrary that j' is pivotal at θ, i.e. $x(\theta) = 1$ and $x(0, \theta_{-j'}) = 0$. If $\hat{\pi}_{j,j'} : \Theta^N \to \Theta^N$ is the function permuting the j-th and j'-th component, then $\hat{\pi}_{j,j'}[(0, \theta_{-j})] \leq (0, \theta_{-j'})$. From monotonicity it follows that $x\left(\hat{\pi}_{j,j'}[(0, \theta_{-j})]\right) = 0$ and symmetry implies that $x(0, \theta_{-j}) = 0$, contradicting the assumption that j is not pivotal at θ.

Step 2: For all i and θ_{-i} we have $h_i(\theta_{-i}) = 0$.

Universal participation immediately implies that an agent with valuation 0 gets a weakly positive utility, i.e., $0 \cdot x(0, \theta_{-i}) + t_i(0, \theta_{-i}) \geq 0$. This implies $h_i(\theta_{-i}) \geq 0$ for all i, θ_{-i}. To obtain a contradiction, suppose that there exists an agent j and a report profile $\theta_{-j} \in \Theta_{-j}$ such that $h_j(\theta_{-j}) > 0$. By step one, we can choose θ_j such that no agent is pivotal at $\theta := (\theta_j, \theta_{-j})$, implying by (1) that $\sum_i t_i(\theta) = \sum_i h_i(\theta_{-i}) > 0$, contradicting (F). □

The following lemma shows how utilitarian welfare of a social choice function can be expressed as the sum of two terms. The first only depends on the allocation rule, and the second consists of the redistribution payments.

Lemma 3. *Let (x, t) be an incentive compatible direct mechanism for social choice rule $G = (X^G, T^G)$ and define*

$$\psi(\theta_i) = \begin{cases} \frac{-F(\theta_i)}{f(\theta_i)} & \text{if } \theta_i \leq 0, \\ \frac{1-F(\theta_i)}{f(\theta_i)} & \text{otherwise.} \end{cases} \tag{2}$$

Then we have

$$U(X^G, T^G) = \int_{\Theta^N} \left[\sum_{i \in N} \psi(\theta_i)\right] x(\theta) dF^N(\theta) + \sum_{i \in N} \int_{\Theta^{N-1}} h_i(\theta_{-i}) dF^{N-1}(\theta_{-i}).$$

Proof. Note that for all θ_{-i},

$$\begin{aligned} &\int_{\underline{\theta}}^{\overline{\theta}} \left[\int_0^{\theta_i} x(\beta, \theta_{-i}) d\beta\right] f(\theta_i) d\theta_i \\ &= \left[\int_0^{\overline{\theta}} x(\beta, \theta_{-i}) d\beta \underbrace{F(\overline{\theta})}_{=1} - \int_0^{\underline{\theta}} x(\beta, \theta_{-i}) d\beta \underbrace{F(\underline{\theta})}_{=0}\right] - \int_{\underline{\theta}}^{\overline{\theta}} x(\theta_i, \theta_{-i}) F(\theta_i) d\theta_i \\ &= \int_0^{\overline{\theta}} \frac{1 - F(\theta_i)}{f(\theta_i)} x(\theta_i, \theta_{-i}) dF(\theta_i) + \int_{\underline{\theta}}^0 \frac{-F(\theta_i)}{f(\theta_i)} x(\theta_i, \theta_{-i}) dF(\theta_i) \end{aligned}$$

$$= \int_{\underline{\theta}}^{\overline{\theta}} \psi(\theta_i)\ x(\theta_i, \theta_{-i}) dF(\theta_i), \tag{3}$$

where the first equality follows from integrating by parts, the second from rearranging terms and the third from the definition of Ψ.

Now rewrite

$$\begin{aligned} U(X^G, T^G) &= \int_{\Theta^N} \sum_{i \in N} \big[\theta_i x(\theta) + t_i(\theta)\big] dF^N(\theta) \\ &= \sum_{i \in N} \int_{\Theta^{N-1}} \int_{\underline{\theta}}^{\overline{\theta}} \left[\int_0^{\theta_i} x(\beta, \theta_{-i}) d\beta + h_i(\theta_{-i}) \right] dF(\theta_i) dF^{N-1}(\theta_{-i}) \\ &= \int_{\Theta^N} \left[\sum_{i \in N} \psi(\theta_i) \right] x(\theta) dF^N(\theta) + \sum_{i \in N} \int_{\Theta^{N-1}} h_i(\theta_{-i}) dF^{N-1}(\theta_{-i}), \end{aligned}$$

where the first equality follows by definition, the second from Lemma 1 and the third by plugging in equation (3). □

For any subset $S \subseteq N$ of the agents, define the corresponding *orthant* as $\mathcal{O}_S = \{\theta \in \Theta^N \mid \theta_i \geq 0 \text{ if } i \in S, \theta_i \leq 0 \text{ if } i \notin S\}$.

Lemma 4. *Suppose that $\psi(\theta)$ is non-increasing in θ and $\int \psi(\theta) dF^N(\theta) < \infty$. Let $\mathcal{O}_S$ be the orthant corresponding to some subset of agents S. Then the problem*

$$\begin{aligned} \max_x \ & \int_{\mathcal{O}_S} \psi(\theta) \cdot x(\theta) dF^N(\theta) \\ \text{s. t. } & x \text{ is non-decreasing in } \theta \\ & 0 \leq x(\theta) \leq 1 \end{aligned}$$

is solved optimally either by setting $x^*(\theta) = 1$ or $x^*(\theta) = 0$.

Proof. Suppose to the contrary that there exists a function $\hat{x}(\theta)$ that achieves a strictly higher value. Let $a_i := \inf\{\theta_i \mid (\theta_i, 0_{-i}) \in \mathcal{O}_S\}$, $b_i := \sup\{\theta_i \mid (\theta_i, 0_{-i}) \in \mathcal{O}_S\}$ and define $x^{(1)}(\theta_1, \theta_{-1}) := \frac{1}{F(b_1) - F(a_1)} \int_{a_1}^{b_1} \hat{x}(\beta, \theta_{-1}) dF(\beta)$. This function is constant in θ_1, feasible for the above problem given that $\hat{x}$ is feasible and, by Chebyshev's inequality, for all θ_{-1},

$$\begin{aligned} &\int_{a_1}^{b_1} \psi(\theta_1, \theta_{-1}) \hat{x}(\theta_1, \theta_{-1}) dF(\theta_1) \\ &\qquad \leq \int_{a_1}^{b_1} \psi(\theta_1, \theta_{-1}) dF(\theta_1) \frac{1}{F(b_1) - F(a_1)} \int_{a_1}^{b_1} \hat{x}(\theta_1, \theta_{-1}) dF(\theta_1) \\ &\qquad = \int_{a_1}^{b_1} \psi(\theta_1, \theta_{-1}) x^{(1)}(\theta_1, \theta_{-1}) dF(\theta_1). \end{aligned}$$

Since this inequality holds point-wise, we also have

$$\int_{\mathcal{O}_S} \psi(\theta)\hat{x}(\theta)dF^N(\theta) \le \int_{\mathcal{O}_S} \psi(\theta)x^{(1)}(\theta)dF^N(\theta).$$

Iteratively defining $x^{(j)}(\theta_j, \theta_{-j}) = \frac{1}{F(b_j)-F(a_j)} \int_{a_j}^{b_j} x^{(j-1)}(\beta, \theta_{-j})dF(\beta)$ for $j = 2, \dots, N$ we get a function $x^{(N)}(\theta)$ that is constant in θ. Repeatedly applying Chebyshev's inequality along every dimension, we get

$$\int_{\mathcal{O}_S} \psi(\theta)\hat{x}(\theta)dF^N(\theta) \le \int_{\mathcal{O}_S} \psi(\theta)x^{(N)}(\theta)dF^N(\theta).$$

Since the objective function is linear in x, the constant function $x^{(N)}$ is weakly dominated by either $x^* \equiv 1$ or $x^* \equiv 0$, contradicting the initial claim. □

Proof of Theorem 1. Lemma 2 and Lemma 3 together imply that for any anonymous social choice function $G = (X^G, T^G)$ it holds that

$$U(X^G, T^G) = \int_{\Theta^N} \left[\sum_{i \in N} \psi(\theta_i) \right] x(\theta)dF^N(\theta),$$

where ψ is defined in (2) and x is the decision rule of the corresponding strategy-proof direct revelation mechanism. Lemma 4 then implies that the optimal allocation rule is constant and equal to 0 or 1 in each orthant. Symmetry of the problem implies that the optimal choice depends only on the number of agents with positive types.

Hence, it remains to determine the optimal cutoff for qualified majority voting. Let k solve

$$k\mathbb{E}[\theta_i \mid \theta_i \ge 0] + (N-k)\mathbb{E}[\theta_i \mid \theta_i \le 0] = 0.$$

Then the expected aggregate valuation, given that $k' < k$ agents are in favor of the proposal, is negative. Therefore, it is optimal to accept the proposal if and only if at least $\lceil k \rceil$ agents have a positive valuation. □

Proof of Corollary 1. Lemma 2 implies that for any social choice function satisfying the requirements of the corollary, one cannot redistribute money back to the agents. Lemma 1 then implies that any budget balanced social choice function must be constant in each orthant. Monotonicity and anonymity then imply that these social choice functions can be implemented by qualified majority voting. □

Chapter 2

Preference Intensities in Repeated Collective Decision-Making

We study welfare-optimal decision rules for committees that repeatedly take a binary decision. Committee members are privately informed about their payoffs and monetary transfers are not feasible. In static environments, the only strategy-proof mechanisms are voting rules which are inefficient as they do not condition on preference intensities. The dynamic structure of repeated decision-making allows for richer decision rules that overcome this inefficiency. Nonetheless, we show that often simple voting is optimal for two-person committees. This holds for many prior type distributions and irrespective of the agents' patience.

1. Introduction

Simple voting rules are known to be inefficient when a majority with weak preferences outvotes a minority with strong preferences. For instance, if ten out of one hundred citizens of a village are willing to pay $20 for changing a law, but the rest has a willingness-to-pay of $1 for keeping the old one, votes would be 90 to 10 against the new law, although it would be efficient to pass it.

Money could be used as a tool to elicit preference intensities and thereby to implement the efficient allocation, but in many situations there are moral or other considerations that prevent the use of monetary means. Instead, this chapter examines the possibilities of using the dynamic structure of environments where group decisions have to be made repeatedly in order to provide incentives for truthful preference revelation. In fact, repeated decision problems are ubiquitous in everyday life, ranging from examples in parliament to hiring committees. In these environments, it is sensible to assume that agents will not proceed myopically from period to period and therefore will not vote sincerely. As Buchanan and Tullock (1962) emphasize, "any rule must be analyzed in terms of the results it will produce, not on a single issue, but on the whole set of issues." Consequently, it is not only reasonable to look at equilibrium behavior under a specific decision rule, but to search for rules that maximize a given objective like, for example, the welfare of the agents.

Consider the following example, which illustrates the possibility of increasing sensitivity to preference intensities: Assume that the decision rule prescribes to accept if at least one of two agents is in favor of the project, unless the other agent uses one of his limited possibilities to exercise a veto. In this situation, agents are faced with a trade-off between the current and future periods. If an agent exercises a veto now, the decision rule decides in her favor, but at the cost of fewer possibilities to use a veto in the future, which reduces the agent's continuation value. Intuitively, agents will use their veto right only if their preference against the proposed project exceeds some threshold. This has the effect that more refined information about the agents' preferences is elicited and potentially a more efficient allocation can be implemented.

Given these ideas, the question is why we see so many decision rules that use simple majority voting in every period, and, more generally, which decision rule is the best in terms of providing the highest welfare to the agents. In this chapter, we tackle the latter question and show that, surprisingly, voting rules are optimal among many reasonable decision rules. This provides a hint to the answer for the former question on why voting is used so universally.

More specifically, we analyze a model with two agents who are repeatedly presented a proposal that they need to either accept or reject. Each agent has a positive or negative willingness-to-pay for accepting the proposal, which is private information and drawn from a distribution function. Due to the revelation principle, we focus on direct mechanisms that simply map past preferences and decisions, and preferences in the current period, into a probability of accepting the current proposal. This allows for the modeling of many conceivable decision rules. We require that decision rules be incentive compatible, so that reporting preferences truthfully is a *periodic ex-post equilibrium*. This means that in any period, given any history, it is a dominant strategy to report the preference truthfully. This requirement renders incentives robust to uncontrolled changes in the information structure as well as deviations of the other player.

We provide a characterization of incentive compatible decision rules in terms of the allocation in a given period and the continuation values the rule promises. Viewing the continuation values as a substitute for money enables us to treat any given decision rule as a static mechanism which can then be improved upon while preserving incentives. The new continuation values of the improved static mechanism can then be implemented by specifying a new dynamic decision rule. As a result, we are able to show that if the preference distributions satisfy an increasing hazard rate condition, then voting rules are optimal within two classes of mechanisms. First, they are optimal among decision rules that satisfy *unanimity*, i.e., rules that never contradict the decision that both agents would unanimously agree on. This is a reasonable robustness requirement since one could expect that the agents will not adhere to the decision rule if they unanimously agree to do something else. Second, if the type distributions are *neutral across alternatives*, i.e., the density is symmetric around zero, then voting rules are also optimal among all deterministic decision rules.

Therefore, if the type distributions are neutral across alternatives, we get the summarizing result that any decision rule yielding higher welfare than every voting rule has both weaknesses of not satisfying unanimity and not being deterministic. This provides a strong rationale for the use of voting rules in the setting we consider and also provides hints on why rules other than voting are not considered in settings with more agents either.

Relation to the Literature

We build upon literature studying decision rules for dynamic settings. Buchanan and Tullock (1962, page 125) note that

> much of the traditional discussion about the operation of voting rules seems to have been based on the implicit assumption that the positive and negative preferences of voters for and against alternatives of collective choice are of approximately equal intensities. Only on an assumption such as this can the failure to introduce a more careful analysis of vote-trading through logrolling be explained.

Buchanan and Tullock (1962) proceed to analyze vote trading. They argue that agents can benefit if they trade their vote on a decision for which they have a weak preference intensity, and in turn get a vote for a future decision. However, it has early been noted that a trade in votes, while being beneficial for the agents involved, might actually reduce aggregate welfare of the whole committee, a fact sometimes called “the paradox of vote trading” (Riker and Brams 1973). A formal analysis of vote trading has been missing until recently, when Casella, Llorente-Saguer and Palfrey (2012) examined in a competitive equilibrium spirit a model of vote trading. They show that vote trading can actually increase welfare in small committees, but is certain to reduce welfare for committees that are large enough.

Instead of relying on agents playing an equilibrium with non-sincere voting so that they can express their preference intensities, one can design specific decision rules that explicitly take intensities into account. Casella (2005) is the first to take this approach in a dynamic setting, in which agents repeatedly decide on a binary choice. He proposes the concept of storable votes: In each period, each agent receives an additional vote and can use some of his votes for the current decision or, alternatively, he can store his additional vote for future usage. By shifting their votes inter-temporally, agents can concentrate their votes on decisions for which they have a strong preference. Casella (2005) shows that this procedure increases welfare of the committee if there are two members and conjectures that in many circumstances this also holds for larger committees. Hortala-Vallve (2012) analyzes a similar proposal for a static setting (meaning that agents are completely informed about their preferences in all decision problems when making the first decision), in which agents face a number of binary decisions.

Going one step further, one can systematically look for the “best” decision rule. Jackson and Sonnenschein (2007) take a mechanism design approach and show that for a static setting the efficient outcome can be approximated even in the absence of money, by linking a large number of independent copies of the decision problem. This result extends to dynamic settings, as long as individuals are arbitrarily patient. This surprising result hinges critically on a number of strong assumptions: each decision problem has to be an identical copy, the designer is required to have the correct prior belief, agents need to be arbitrarily patient and their beliefs about other agents have to be identical to the common prior. In an attempt to find more robust decision rules, Hortala-Vallve (2010) characterizes the set of strategy-proof decision rules for a static problem. Given that strategy-proofness is a strong requirement in multi-dimensional settings, it is not too surprising that voting rules are the only decision rules that satisfy this restriction.

In contrast, our focus on periodic ex-post equilibrium implies that on the one hand, the set of implementable decision rules is very rich, but on the other hand our results are robust and the optimal mechanism is bounded away from attaining the first-best.

The chapter is structured as follows: In Section 2 we present our model in detail. The results are presented in Section 3 and discussed in Section 4. Some proofs are omitted from the main text and relegated to the appendix.

2. Model

There are two agents who are repeatedly faced with a proposal and have to accept or reject each proposal. Periods are indexed by $t = 0, 1, ... \in T = \mathbb{N}$. The type of an agent i in a given period t is denoted by θ_{it} and indicates his willingness-to-pay for the proposal. Type spaces and distribution functions are the same for each period and each agent, denoted by Θ_i and F respectively, and types are drawn independently across time and agents. We denote by $\tilde{\theta}_{it}$ the random variable corresponding to the type of agent i, and by θ_t a type profile which is an element of the product type space Θ.

In each period, a decision $x_t \in \{0, 1\}$ has to be made. We denote the sequence of decisions up to period t by x^t, and similarly for a sequence of types θ_i^t. Accordingly, for an infinite sequence we write x^T.

Mechanisms

In this model a dynamic version of the revelation principle holds (Myerson (1986), for similar arguments see Pavan, Segal and Toikka (2008)), hence we can focus on truthfully implementable direct revelation mechanisms.

Definition 1. *A* mechanism χ *is a sequence of decision rules* $\{\chi_t\}_{t \in T}$ *that map past decisions and type profiles into a distribution over decisions in the current period:*

$$\chi_t : \Theta^t \times \{0, 1\}^{t-1} \to [0, 1].$$

Preferences

Agents have linear von-Neumann-Morgenstern utility functions and there are no monetary payments. Given a period t and a decision x_t for this period, the utility of agent i with type θ_{it} is $v_{it}(\theta_{it}, x_t) = \theta_{it} x_t$. Agents discount the future with the common discount factor $\delta \in [0, 1)$. Consequently, utility of agent i with type sequence θ_i^T is

$$V_i(\theta_i^T, x^T) = \sum_{t \in T} \delta^t \theta_{it} x_t$$

for the decision sequence x^T.

Equilibrium Concept and Incentive Compatibility

In every period t, agent i learns about his preference type θ_{it}, which is his private information, and then sends a report r_{it}. The history known to the designer in period t, $h^t = (x^{t-1}, r^{t-1})$, consists of past decisions and past reports.

Given a mechanism χ, we can write the value function for agent i:

$$W_i(h^t, \theta_t) = \sup_{r_{it} \in \Theta_i} \theta_{it} \chi_t(h^t, r_{it}, \theta_{-it}) + \delta \mathbb{E}_{\Theta_{t+1}} W_i(h^{t+1}, \tilde{\theta}_{t+1}) \tag{1}$$

Here, h^{t+1} is the history in the next period, consisting of $\chi_t(h^t, r_{it}, \theta_{-it})$ and (r_{it}, θ_{-it}) appended to h^t. The valuation function specifies, given any history h^t, and the current type profile θ_t, the highest utility the agent can possibly obtain for some report r_{it}, assuming that she reports optimally in the future and the other agents report truthfully. Given a specific history h^t, the mechanism χ induces an allocation rule and continuation functions which we will denote

$$x_t(\theta_t) = \chi_t(h^t, \theta_t) \quad \text{and}$$
$$w_{it}(\theta_t) = \delta \mathbb{E}_{\Theta_{t+1}} W_i(h^{t+1}, \tilde{\theta}_{t+1}).$$

If the current period is clear from the context, we will also drop the subscript t. The pair (x_t, w_t) is called the *stage mechanism after history* h_t and we say that w_t is *generated* by the mechanism χ. A stage mechanism is *admissible* if it is generated by some mechanism χ.

Definition 2. *A mechanism is* periodic ex-post incentive compatible (IC) *if for every period t and for all histories h^t the following holds: For every θ_{-i} and every θ_i we have that*

$$\theta_{it} x(\theta_{it}, \theta_{-it}) + w_{it}(\theta_{it}, \theta_{-it}) \geq \theta_{it} x(r_{it}, \theta_{-it}) + w_{it}(r_{it}, \theta_{-it}) \tag{2}$$

for all reports $r_i \in \Theta_i$.

See, e.g., Athey and Miller (2007), Bergemann and Välimäki (2010). The definition in particular states that if a mechanism is incentive compatible, then every stage mechanism for all histories is incentive compatible. The following lemma can be proved using the Envelope Theorem (which is a standard exercise in mechanism design).

Lemma 1. *A mechanism is IC if and only if for each agent i the following two conditions hold:*

1. *Monotonicity of x: $x(\theta_i, \theta_{-i}) \leq x(\theta_i', \theta_{-i})$ for $\theta_i \leq \theta_i'$.*
2. *Payoff equivalence: Fix $\hat{\theta}_i \in \Theta_i$. Then for all θ*

$$\theta_i x(\theta_i, \theta_{-i}) + w_i(\theta_i, \theta_{-i}) = \hat{\theta}_i x(\hat{\theta}_i, \theta_{-i}) + w_i(\hat{\theta}_i, \theta_{-i}) + \int_{\hat{\theta}_i}^{\theta_i} x(\beta, \theta_{-i}) d\beta. \tag{3}$$

Since the term $\hat{\theta}_i x(\hat{\theta}_i, \theta_{-i}) + w_i(\hat{\theta}_i, \theta_{-i})$ is independent of θ_i, we will write $h_i(\theta_{-i})$ for it. Note, however, that $h_i(\theta_{-i})$ does depend on the particular choice of $\hat{\theta}_i$.

Objective

For a given stage mechanism we can write down the expected welfare going forward from period t as

$$U_{h^t}(\chi) = U_{h^t}(x, w) := \mathbb{E}_{\Theta_t}\big[(\theta_1 + \theta_2)x_t(\theta) + w_{1t}(\theta) + w_{2t}(\theta)\big].$$

This is the period-t expected discounted welfare that the agents receive after history h^t. The aim of this chapter is to identify welfare-optimal mechanisms, that is, mechanisms χ that solve

$$\max_{\chi} \ U(\chi) := U_{h^0}(\chi), \quad \text{s. t.} \quad \chi \text{ is IC.}$$

Lemma 2 in the appendix provides a useful way to rewrite the objective function in terms of the allocation rule x and $h_i(\theta_{-i})$.

3. Results

The aim of this section is to identify mechanisms that are optimal in the above stated sense. The following conditions on F which we need to derive our results are standard in the mechanism design literature.

Condition 1 (Monotone Hazard Rates)**.** *The hazard rate* $\frac{f(\theta_i)}{1-F(\theta_i)}$ *is non-decreasing in* θ_i *and the reversed hazard rate* $\frac{f(\theta_i)}{F(\theta_i)}$ *is non-increasing in* θ_i*.*

A *voting rule* x is a rule where $x(\theta)$ only depends on $\{sgn(\theta_i)\}_{i=1,2}$. A voting mechanism is a mechanism where the allocation rule after all histories is a voting rule. In each of the two subsections below we will present a setting in which the welfare-maximizing dynamic decision rule is a voting mechanism.

The proofs in each part will proceed as follows: First, we show that under the appropriate assumptions stage mechanisms consisting of a voting rule and promising the same continuation payoffs for all type profiles are weakly welfare-superior to all other stage mechanisms. Then we make use of the following proposition to deduce that also the best dynamic mechanism uses a voting rule in every period. For this step to work it is helpful that optimal stage mechanisms are of as simple a form as voting mechanisms.

Proposition 1. *Assume that for every history* h^t *and admissible stage mechanism* (x_t, w_t) *in period* t*, there exists an admissible stage mechanism* $(\hat{x}_t, \hat{w}_t)$*, where* $\hat{x}_t$ *is a voting rule and* $\hat{w}_t$ *is constant, and such that*

$$U_{h^t}(x_t, w_t) \leq U_{h^t}(\hat{x}_t, \hat{w}_t).$$

Then a voting mechanism is among the optimal mechanisms.

Proof. We start with any dynamic mechanism χ and transform it into a mechanism that uses a voting rule in every period and such that U weakly increases. Start with $t = 0$. The assumption states that there exists a voting stage mechanism $(\hat{x}_0, \hat{w}_0)$ with constant $\hat{w}_0$ and such that $U(\hat{x}_0, \hat{w}_0) \geq U(x_0, w_0)$. Since the voting stage mechanism is admissible and promises constant continuations, these continuations can be generated by a mechanism that is independent of h^1. Denote by χ' this new dynamic mechanism. Since x_1' and w_1' are independent of h^1, we know (again by the assumption) that there exists a voting stage mechanism $(\hat{x}_1, \hat{w}_1)$ with constant $\hat{w}_1$ and such that $U_{h^1}(\hat{x}_1, \hat{w}_1) \geq U_{h^1}(x_1', w_1')$ for all h^1. Again, $\hat{w}_1$ can be generated by a mechanism that does not condition on histories h^2. Now if we let χ'' be the mechanism that arises if one exchanges the stage mechanism (x_1', w_1') in χ' for $(\hat{x}_1, \hat{w}_1)$, we know that χ'' is still incentive compatible: All promised continuations in period 0 change by the same amount, independent of the history h^1 and in particular independent of θ_0. Repeating this argument inductively for $t \geq 2$ completes the proof. □

Unanimity

Unanimity requires the mechanism to always adhere to a decision to which both agents agree. For example, if both types in some period are positive the mechanism has to choose $x_t = 1$ for sure. Formally, the condition is defined as follows:

Definition 3. *A mechanism is called* unanimous *if, for every period and all possible histories,* $x(\theta) = 1$ *if* $\theta > 0$ *and* $x(\theta) = 0$ *if* $\theta < 0$.

Note that mechanisms not satisfying this requirement will probably have legitimacy problems: Although all parties involved in the decision process opt in favor of the proposal, the mechanism forces its rejection. Furthermore, if agents are not able to collectively commit to the decision prescribed by the mechanism, then mechanisms satisfying unanimity are the only feasible mechanisms. Also note that mechanisms proposed in the literature are not excluded by this assumption (see, e. g., Jackson and Sonnenschein 2007, Casella 2005). In the next subsection we will see that even when relaxing this restriction, for certain distribution functions only non-deterministic decision rules can yield a higher expected welfare than voting rules.

Theorem 1. *Suppose that* F *satisfies Condition 1. Then a voting mechanism is optimal among all unanimous mechanisms.*

Proof. The proof consists of establishing the preconditions of Proposition 1. So let (x, w) be a stage mechanism after some history h^t (since we are only concerned with unanimous mechanisms, x satisfies unanimity). Set $(\hat{\theta}_1, \hat{\theta}_2) = (0, 0)$ and let h_i be the resulting redistribution functions implied by Lemma 1. Let $\theta^* \in \arg\max_{\theta \in \Theta_i} h_1(\theta) + h_2(\theta)$. We first show that setting $h_1(\theta_2) = h_1(\theta^*)$ for all θ_2 and $h_2(\theta_1) = h_2(\theta^*)$ for all θ_1 does not decrease $U_{h^t}(x, w)$.

Since so far we have not changed x, by Lemma 2 it is enough to show that the terms involving the redistribution functions do not decrease in this step. But this follows from

$$\begin{aligned}\int_{\Theta_1} h_2(\theta_1) dF(\theta_1) + \int_{\Theta_2} h_1(\theta_2) dF(\theta_2) &= \int_{\underline{\theta}}^{\overline{\theta}} \big[h_2(\beta) + h_1(\beta)\big] dF(\beta) \\ &\leq \int_{\underline{\theta}}^{\overline{\theta}} \big[h_2(\theta^*) + h_1(\theta^*)\big] dF(\theta^*).\end{aligned}$$

Next we show that changing x to a voting rule does not decrease welfare. It is enough to consider the regions where $\theta_1 \leq 0, \theta_2 \geq 0$ and $\theta_1 \geq 0, \theta_2 \leq 0$ because the mechanism is unanimous. By Lemma 3 and the choice of $(\hat{\theta}_1, \hat{\theta}_2)$, we know that the first term in (4), which for the region $\theta_1 \leq 0, \theta_2 \geq 0$ amounts to

$$\int_{\underline{\theta}}^{0} \int_{0}^{\bar{\theta}} \left[\frac{-F(\theta_1)}{f(\theta_1)} + \frac{1 - F(\theta_2)}{f(\theta_2)} \right] x(\theta_1, \theta_2) dF(\theta_2) dF(\theta_1),$$

is maximized by setting x to 1, as soon as Condition 1 holds. Since the same is true for the region where $\theta_1 \geq 0, \theta_2 \leq 0$, we have constructed a voting stage mechanism that is weakly welfare superior to the old stage mechanism.

Let (x', w') denote the new stage mechanism. The proof is complete if we can show that w' is constant and can be generated. Constancy of w' holds for any stage mechanism where x' is a voting rule and the functions h'_i are constant. More specifically, w'_i is equal to $h_i(\theta^*)$. Since the old mechanism was unanimous, $w_i(\theta^*, \theta^*) = h_i(\theta^*)$. Because $w_i(\theta^*, \theta^*)$ could be generated, it follows that w' can be generated. □

Neutrality of Alternatives

In this section, we show that in some situations we can derive optimality of voting mechanisms even if unanimity does not hold. This shows that the restriction imposed in the previous section does in many cases not reduce welfare.

We assume that the distribution of types is *neutral across alternatives*, i.e., it is symmetric around 0. This is an important special case of our general model and has been analyzed, among others, by Carrasco and Fuchs (2011). For instance, this assumption is satisfied if a committee has to decide among two proposals that are valued equally ex ante. Specifying one alternative as the default, the distribution of valuations for changing from the default to the alternative proposal is symmetric around 0.

Theorem 2. *Suppose F satisfies Condition 1 and is neutral across alternatives. Then a voting mechanism is optimal among all deterministic mechanisms.*

The proof of Theorem 2 is presented in the appendix. Similar arguments as in the last subsection can be given for restricting attention to deterministic mechanisms: First, stochastic mechanisms are difficult to implement and face legitimacy problems in practice. It is barely conceivable that a parliament would introduce decision protocols that involve random elements. Second, all proposed mechanisms in the literature and mechanisms observed in practice are usually deterministic and therefore not excluded from our analysis. Numerical simulation also suggests that expected welfare can be improved only slightly using stochastic mechanisms. The following corollary combines Theorem 1 and Theorem 2 and summarizes all properties one has to give up in order to improve upon voting rules.

Corollary 1. *Assume F satisfies Condition 1 and is neutral across alternatives. Then every decision rule that is strictly welfare-superior to any voting rule is stochastic and does not satisfy unanimity.*

4. Discussion

We have seen that despite the absence of money as a means for implementing rules other than majority voting, the possibility to condition decision rules on the past gives us the possibility to design dynamic decision rules that take preference intensities into account. However, we have shown that for committees consisting of two players the welfare maximizing dynamic decision rule nonetheless consists of simple majority voting in every period. This holds unless desirable properties of the decision rule are given up. We therefore provide a possible explanation for why majority voting is used almost universally in practice.

One extension of our model is to allow for correlation of agent types over time. However, this restricts the class of incentive compatible mechanisms since the quasi-linear separation of continuation payoffs from the payoff in the current period disappears. While voting rules would still be optimal in this restricted class, our model without correlation shows that voting rules are also optimal in the larger class.

A major open problem is the question as to what extent our results generalize to more than two agents. We believe that a substantial difficulty towards progress in this direction is to understand in how far continuation values can be redistributed among the agents.

Appendix

Helpful Lemmata

The following shows how the welfare of every incentive compatible mechanism can be expressed in terms of the allocation function and the functions h_i defined following Lemma 1.

Lemma 2. *Let χ be an incentive compatible mechanism and define*

$$\psi(\theta_i) = \begin{cases} \frac{-F(\theta_i)}{f(\theta_i)} & \text{if } \theta_i \leq \hat{\theta}_i, \\ \frac{1-F(\theta_i)}{f(\theta_i)} & \text{otherwise.} \end{cases}$$

Then for every history h^t we have

$$U_{h^t}(\chi) = \int_{\Theta} \big[\psi(\theta_1) + \psi(\theta_2)\big] x(\theta) dF(\theta) + \int_{\Theta_1} h_2(\theta_1) dF(\theta_1) + \int_{\Theta_2} h_1(\theta_2) dF(\theta_2). \quad (4)$$

Proof. First note that

$$U_{h^t}(\chi) = \int_{\underline{\theta}}^{\overline{\theta}} \int_{\underline{\theta}}^{\overline{\theta}} \Big[\theta_1 x(\theta) + \theta_2 x(\theta) + w_1(\theta) + w_2(\theta)\Big] dF(\theta_2) dF(\theta_1), \quad (5)$$

and by Lemma 1

$$w_i(\theta) = \int_{\hat{\theta}_i}^{\theta_i} x(\beta, \theta_{-i}) d\beta - \theta_i x(\theta) + h_i(\theta_{-i}). \quad (6)$$

Using integration by parts, we first rewrite the term

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[\int_{\hat{\theta}_i}^{\theta_i} x(\beta, \theta_{-i}) d\beta\right] f(\theta_i) d\theta_i$$
$$= \left[\int_{\hat{\theta}_i}^{\overline{\theta}} x(\beta, \theta_{-i}) d\beta \underbrace{F(\overline{\theta})}_{=1} - \int_{\hat{\theta}_i}^{\underline{\theta}} x(\beta, \theta_{-i}) d\beta \underbrace{F(\underline{\theta})}_{=0}\right] - \int_{\underline{\theta}}^{\overline{\theta}} x(\theta_i, \theta_{-i}) F(\theta_i) d\theta_i$$
$$= \int_{\hat{\theta}_i}^{\overline{\theta}} \frac{1 - F(\theta_i)}{f(\theta_i)} x(\theta) dF(\theta_i) + \int_{\underline{\theta}}^{\hat{\theta}_i} \frac{-F(\theta_i)}{f(\theta_i)} x(\theta) dF(\theta_i). \tag{7}$$

Now plug (6) into (5) and use (7) to complete the proof. □

The next lemma implies, together with Condition 1, that the first part of (4) is maximized by a constant allocation function whenever only one part of the function ψ is considered.

Lemma 3. *Suppose that $\psi(\theta_1, \theta_2)$ is non-increasing in θ_1 and θ_2, and that $\int \psi(\theta) dF(\theta) < \infty$. Then the problem*

$$\max_x \int_a^b \int_c^d \psi(\theta_1, \theta_2) \cdot x(\theta_1, \theta_2) dF(\theta_2) dF(\theta_1)$$
$$\text{s.t. } x \textit{ is non-decreasing in } \theta$$
$$0 \le x(\theta) \le 1$$

is solved optimally either by setting $x^*(\theta) = 1$ or $x^*(\theta) = 0$.

Proof. Suppose to the contrary that there exists a function $\hat{x}(\theta)$ that achieves a strictly higher value. Define $x'(\theta_1, \theta_2) := \frac{1}{F(d)-F(c)} \int_c^d \hat{x}(\theta_1, \beta) dF(\beta)$. This function is feasible for the above problem given that $\hat{x}$ is feasible and, by Chebyshev's inequality, for all θ_1,

$$\int_c^d \psi(\theta_1, \theta_2) \hat{x}(\theta_1, \theta_2) dF(\theta_2)$$
$$\le \int_c^d \psi(\theta_1, \theta_2) dF(\theta_2) \frac{1}{F(d) - F(c)} \int_c^d \hat{x}(\theta_1, \theta_2) dF(\theta_2)$$
$$= \int_c^d \psi(\theta_1, \theta_2) x'(\theta_1, \theta_2) dF(\theta_2).$$

Since this inequality holds for every θ_1, we also have

$$\int_a^b \int_c^d \psi(\theta_1, \theta_2) \hat{x}(\theta_1, \theta_2) dF(\theta_2) dF(\theta_1) \le \int_a^b \int_c^d \psi(\theta_1, \theta_2) x'(\theta_1, \theta_2) dF(\theta_2) dF(\theta_1).$$

Defining $x''(\theta_1, \theta_2) = \frac{1}{F(b)-F(a)} \int_a^b x'(\theta_1, \theta_2) dF(\theta_1)$ and again applying Chebyshev's inequality as above, we get that

$$\int_a^b \int_c^d \psi(\theta_1, \theta_2) x'(\theta_1, \theta_2) dF(\theta_2) dF(\theta_1) \le \int_a^b \int_c^d \psi(\theta_1, \theta_2) x''(\theta_1, \theta_2) dF(\theta_2) dF(\theta_1).$$

Since the objective function is linear in x, the constant function x'' is weakly dominated by either $x \equiv 1$ or $x \equiv 0$, contradicting the initial claim. $\square$

Proof of Theorem 2

Proof. We establish the preconditions of Proposition 1. Fix an arbitrary history h_t and consider the stage mechanism (x, w) employed after this history. Let $\overline{w} := max_\theta\{w_1(\theta) + w_2(\theta)\}$ and let θ_w be an optimizer. We normalize w such that $w_1(\theta_w) = w_2(\theta_w) = 0$ by decreasing w_i by $w_i(\theta_w)$ for all i. This does not affect incentive compatibility. After the normalization we have

$$w_1(\theta) + w_2(\theta) \leq 0.$$

We start with some preliminaries where we derive a set of inequalities that are satisfied by every incentive compatible stage mechanism for which the above inequality holds.

Preliminaries:

Set $(\hat{\theta}_1, \hat{\theta}_2) := (\overline{\theta}, \underline{\theta})$, let h_i denote the resulting redistribution functions implied by Lemma 1 and define $g_i(\theta) := \theta_i x(\theta) - \int_{\hat{\theta}_i}^{\theta_i} x(\beta, \theta_{-i}) d\beta$. It follows from Lemma 1 that $w_i(\theta) = -g_i(\theta) + h_i(\theta_{-i})$. Let $h^* := \max_\theta\{h_1(\theta) + h_2(-\theta)\} - \overline{\theta}$ and θ^* be a maximizer. Normalize h such that $h_1(\theta^*) = h^* + \overline{\theta}$ and $h_2(-\theta^*) = 0$ by increasing $h_1(x_2)$ and decreasing $h_2(x_1)$ by $h_2(-\theta^*)$. The definition of h^* implies

$$h_1(\theta) + h_2(-\theta) \leq h^* + \overline{\theta} \qquad \text{for all } \theta, \tag{8}$$

and $w_1(\theta, -\theta) + w_2(\theta, -\theta) \leq 0$ implies

$$\begin{aligned} h_1(\theta) + h_2(-\theta) &\leq g_1(-\theta, \theta) + g_2(-\theta, \theta) \\ &= -\int_{\overline{\theta}}^{-\theta} x(\beta, \theta) d\beta - \int_{\underline{\theta}}^{\theta} x(-\theta, \beta) d\beta \\ &\leq \int_{-\theta}^{\overline{\theta}} x(\beta, \theta) d\beta \leq \overline{\theta} + \theta. \end{aligned} \tag{9}$$

By plugging θ^* into (9) and using the definition of h^*, it follows that $h^* \leq \theta^*$.

Define $a := \inf\{\theta_1 \mid x(\theta_1, h^*) = 1\}$. If there does not exist θ_1 such that $x(\theta_1, h^*) = 1$, set $a := \overline{\theta}$. Without loss we can assume that $a \geq -h^*$, since otherwise we can "mirror" the mechanism on the dotted line shown in Figure 1.[1] Let $\theta_1 \geq a$. Then expanding and rearranging $w_1(\theta_1, \theta^*) + w_2(\theta_1, \theta^*) \leq 0$ yields

$$\begin{aligned} h_2(\theta_1) &\leq -(h^* + \overline{\theta}) + g_1(\theta_1, \theta^*) + g_2(\theta_1, \theta^*) \\ &= -h^* - \overline{\theta} + \theta_1 - \int_{\overline{\theta}}^{\theta_1} x(\beta, \theta^*) d\beta + \theta^* - \int_{\underline{\theta}}^{\theta^*} x(\theta_1, \beta) d\beta \end{aligned}$$

[1] Let $(x^{\#}, w^{\#})$ be the mirrored mechanism, then $x^{\#}(\theta_1, \theta_2) = 1 - x(-\theta_2, -\theta_1)$, $w_i^{\#}(\theta_1, \theta_2) = w_{-i}(-\theta_2, -\theta_1)$. The new mechanism is IC iff. the old mechanism is IC and by our symmetry assumptions the mirrored mechanism yields the same welfare. Also, h^* and θ^* will not be changed by this operation.

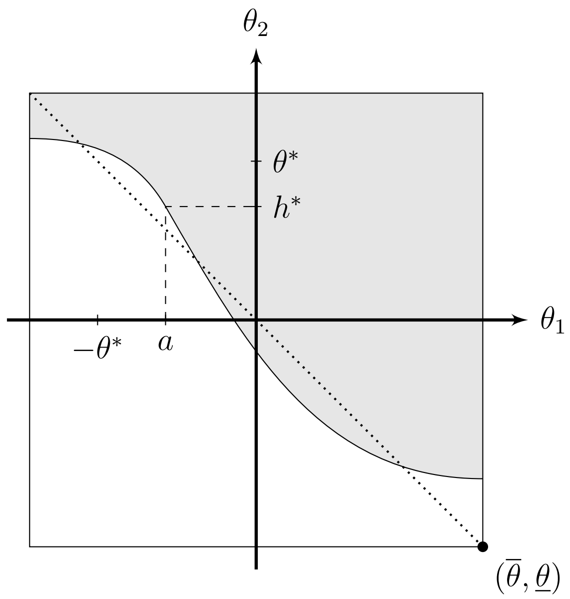

Figure 1: Proof of Theorem 2. The shaded area indicates the profiles θ where $x(\theta) = 1$.

$$= -h^* + \theta^* - \theta^* + h^* - \int_{\underline{\theta}}^{h^*} x(\theta_1, \beta) d\beta$$

$$= - \int_{\underline{\theta}}^{h^*} x(\theta_1, \beta) d\beta, \tag{10}$$

where in the second equality we made use of the fact that $x(\beta, \theta^*) = 1$ for $\beta \geq a$ and $x(\theta_1, \beta) = 1$ for $\theta_1 \geq a$, $h^* \leq \beta \leq \theta^*$ (see Figure 1). Similar arguments will be used more often in the equalities below.

Define $b := \inf\{\theta_2 \mid x(-h^*, \theta_2) = 1\}$ (if there is no θ_2 such that $x(-h^*, \theta_2) = 1$, set $b := \overline{\theta}$) and let $\theta_2 \leq b$. Then $w_1(-\theta^*, \theta_2) + w_2(-\theta^*, \theta_2) \leq 0$ implies

$$h_1(\theta_2) \leq g_1(\theta^*, \theta_2) + g_2(\theta^*, \theta_2)$$

$$= 0 - \int_{\overline{\theta}}^{-\theta^*} x(\beta, \theta_2) d\beta - \int_{\underline{\theta}}^{\theta_2} x(-\theta^*, \beta) d\beta$$

$$= \int_{-\theta^*}^{\overline{\theta}} x(\beta, \theta_2) d\beta. \tag{11}$$

Since by Lemma 1 an incentive compatible stage mechanism is completely determined by x and h, we will in the following change x and h in a number of consecutive steps while making sure that x stays monotone and we never decrease the welfare $U^{h_t}(x, h) := U^{h_t}(x, w)$. At the end of the proof we will make sure that the resulting mechanism is admissible. First, we increase $h_2(\theta_1)$ for $\theta_1 \geq a$ and $h_1(\theta_2)$ for $\theta_2 \leq b$ until (10) and (11) hold with equality since this trivially weakly increases welfare.

Step 1:

In this step we will change the variables $x(\theta)$ with $\theta \in A := \{(\theta_1, \theta_2) \mid \theta_1 \geq a, \theta_2 \leq h^*\}$, $h_2(\theta_1)$ with $\theta_1 \geq a$ and $h_1(\theta_2)$ with $\theta_2 \leq h^*$. If we change h_1 and h_2 such that (11) and (10) continue to hold with equality, we can express changes of all the variables in terms of changes of x. Making use of the fact that for $\theta_2 \leq h^*$, (11) is equivalent to

$$h_1(\theta_2) = \int_a^{\overline{\theta}} x(\beta, \theta_2) d\beta,$$

and by substituting (11) and (10), we can rewrite the the part of U_{h^t} that depends on changes of the variables $x(\theta)$ for $\theta \in A$ as

$$\begin{aligned}
&\int_a^{\overline{\theta}} \int_{\underline{\theta}}^{h^*} \left[\frac{-F(\theta_1)}{f(\theta_1)} + \frac{1 - F(\theta_2)}{f(\theta_2)} \right] x(\theta_1, \theta_2) dF(\theta_2) dF(\theta_1) \\
&\qquad + \int_{\underline{\theta}}^{h^*} \int_a^{\overline{\theta}} x(\beta, \theta_2) d\beta \, dF(\theta_2) - \int_a^{\overline{\theta}} \int_{\underline{\theta}}^{h^*} x(\theta_1, \beta) d\beta \, dF(\theta_1) \\
&= \int_a^{\overline{\theta}} \int_{\underline{\theta}}^{h^*} \left[\frac{1 - F(\theta_1)}{f(\theta_1)} + \frac{-F(\theta_2)}{f(\theta_2)} \right] x(\theta_1, \theta_2) dF(\theta_2) dF(\theta_1).
\end{aligned}$$

Lemma 3 implies that this term is maximized by setting $x(\theta) = 0$ or 1 for $\theta \in A$. To see that we cannot gain by setting $x(\theta) = 1$ we bound

$$\begin{aligned}
U_{h^t}(1, h) &= \int_a^{\overline{\theta}} \int_{\underline{\theta}}^{h^*} \left[\frac{1 - F(\theta_1)}{f(\theta_1)} + \frac{-F(\theta_2)}{f(\theta_2)} \right] dF(\theta_2) dF(\theta_1) \\
&= \int_{\underline{\theta}}^{-a} \int_{\underline{\theta}}^{h^*} \left[\frac{F(\theta_1)}{f(\theta_1)} + \frac{-F(\theta_2)}{f(\theta_2)} \right] dF(\theta_2) dF(\theta_1) \\
&= \int_{\underline{\theta}}^{-a} \int_{-a}^{h^*} \left[\frac{F(\theta_1)}{f(\theta_1)} + \frac{-F(\theta_2)}{f(\theta_2)} \right] dF(\theta_2) dF(\theta_1) \\
&\leq 0 = U_{h^t}(0, h).
\end{aligned}$$

Here, the second equality is due to the symmetry of F around zero, the third equality is because the integral over $[\underline{\theta}, -a] \times [\underline{\theta}, -a]$ vanishes, and the inequality is due to log-concavity of F and the fact that $-a \leq h^*$. Hence, we weakly increase welfare by setting $x \equiv 0$ in A and h_1 and h_2 according to (11) and (10), respectively.

Step 2:

For this step define the set $B = \{\theta_1 > -h^*, \theta_2 > h^* \mid x(\theta_1, \theta_2) = 0\}$. Set $x(\theta) = 1$ for $\theta \in B$ and $h_1(\theta_2) = h^* + \overline{\theta}$ for all θ_2 for which there is a θ_1 such that $(\theta_1, \theta_2) \in B$. We claim that this does not decrease U_{h^t}. Since allocative efficiency improved in this step, we only need to check that the sum of promised continuations increased. First, let $(\theta_1, \theta_2) \in B$. Then (11) is equivalent to

$$h_1(\theta_2) = \int_{-h^*}^{\overline{\theta}} x(\beta, \theta_2) d\beta.$$

Continuations before this change are given by

$$h_2(\theta_1) + h_1(\theta_2) + \int_{\overline{\theta}}^{\theta_1} x(\beta, \theta_2) d\beta = h_2(\theta_1) + \int_{-h^*}^{\overline{\theta}} x(\beta, \theta_2) d\beta + \int_{\overline{\theta}}^{\theta_1} x(\beta, \theta_2) d\beta = h_2(\theta_1).$$

After the change we get

$$h_2(\theta_1) + h^* + \overline{\theta} - \theta_1 + \int_{\overline{\theta}}^{\theta_1} \underbrace{x(\beta, \theta_2)}_{=1} d\beta - \theta_2 + \int_{h^*}^{\theta_2} \underbrace{x(\theta_1, \beta)}_{=1} d\beta = h_2(\theta_1).$$

Fixing $(\theta_1, \theta_2) \in B$, the claim can similarly be shown for points of the form (θ_1', θ_2) and (θ_1, θ_2') where $\theta_2' > \theta_2$.

Step 3:

We claim that setting $x(\theta) = 1$ or $x(\theta) = 0$ for $\theta \in [\underline{\theta}, -h^*] \times [h^*, \overline{\theta}]$ increases U_{h^*}. This follows from the fact that, since, ignoring the part which depends on h_i, the objective function in the area where we change x has the form required by Lemma 3. Symmetry implies that $x(\theta) = 0$ gives the same welfare as $x(\theta) = 1$.

Step 4:

Note that the original mechanism satisfied

$$h_1(-\theta) + h_2(\theta) \leq h^* + \overline{\theta}.$$

Therefore, welfare is not decreased by setting $h_2(\theta) := 0$ and $h_1(-\theta) = h^* + \overline{\theta}$ for $\theta \leq -b$.

Note that the changed mechanism satisfies $w_1(\theta, -\theta) + w_2(\theta, -\theta) \leq 0$: For $a \leq \theta$ this holds as we assumed (10) and (11) to be binding in Step 1, hence $g_1(\theta, -\theta) = g_2(\theta, -\theta) = h_1(-\theta) = h_2(\theta) = 0$. For $-h^* \leq \theta \leq a$, this holds as continuations weren't changed for these values (changed Pivot payments were offset by changes in the h functions, as (11) was assumed to hold with equality in Step 1). For $-b \leq \theta \leq -h^*$ this holds as constraints were assumed to bind in Step 2. For $\underline{\theta} \leq \theta \leq -b$ this holds as $h_1(-\theta) + h_2(\theta) \leq h^* + \overline{\theta} = g_1(\theta, -\theta) + g_2(\theta, -\theta)$.

The fact that $w_1(\theta, -\theta) + w_2(\theta, -\theta) \leq 0$ implies that $h_1(-\theta) + h_2(\theta) \leq g_1(\theta, -\theta) + g_2(\theta, -\theta)$. We can increase h so that equality holds, thereby again improving the mechanism, ending up with the following stage mechanism:

$$x(\theta) = \begin{cases} 1 & \text{if } \theta_2 \geq h^* \\ 0 & \text{else,} \end{cases}$$

$$h_1(\theta_2) = \begin{cases} 0 & \text{if } \theta_2 \leq h^* \\ h^* + \overline{\theta} & \text{else,} \end{cases}$$
$$h_2(\theta_1) = 0.$$

We call this class of mechanisms *phantom dictatorship with parameter* h^*.

Step 5:

So far we have shown that every stage mechanism can be modified until it is a phantom dictatorship while weakly improving welfare. To prove that for every stage mechanism there is a simple voting stage mechanism with weakly higher welfare, we show that simple voting weakly welfare-dominates every phantom dictatorship: Indeed, the optimal phantom dictatorship is given by the parameter $h^* = \mathbb{E}[\theta]$. Therefore, symmetry of F around 0 implies that the optimal phantom dictatorship is characterized by $h^* = 0$, which has the same aggregate welfare as unanimity voting.

The voting stage mechanism we have constructed so far has the continuations profile $w_1(\theta) = w_2(\theta) = 0$ for all θ. It remains to show that this mechanism is admissible. But this follows from the fact that $(0, 0)$ was an implementable continuation profile of the original mechanism (namely, at the type profile θ_w). We therefore established the conditions for Proposition 1, which completes the proof of the theorem. □

Chapter 3

Optimal Private Good Allocation: The Case for a Balanced Budget

In an independent private value auction environment, we are interested in strategy-proof mechanisms that maximize the agents' residual surplus, that is, the utility derived from the physical allocation minus transfers accruing to an external entity. We find that, under the assumption of an increasing hazard rate of type distributions, an optimal deterministic mechanism never extracts any net payments from the agents, that is, it will be budget-balanced. Specifically, optimal mechanisms have a simple "posted price" or "option"form. In the bilateral trade environment, we obtain optimality of posted price mechanisms without any assumption on type distributions.

1. Introduction

Most parts of the mechanism design literature studying welfare maximization problems focus on mechanisms implementing the efficient allocation. However, in general it is not possible to implement the efficient allocation in dominant strategies using budget-balanced mechanisms (Green and Laffont 1979). Given this result, we study the question of how to choose among different mechanisms that cannot attain both, allocative efficiency and budget-balancedness. Since we are concerned with welfare maximization, the social planner's objective function should consist of the agents' aggregate utility and therefore include aggregate transfers. In other words, one seeks to find mechanisms that maximize what we call residual surplus. This is the surplus, or utility, the agents derive from the chosen physical allocation, reduced by the amount of transfers that are lost to an external agency (this is often called "money burning").

A common approach is to implement the efficient allocation via Groves mechanisms and to redistribute as much money to the agents as possible without distorting incentives (Cavallo 2006, Guo and Conitzer 2009, Guo and Conitzer 2010, Moulin 2009). This approach aims at characterizing the optimal mechanism for allocating private goods that implements the *efficient allocation* in *dominant strategies*, is *individually rational* and *never creates a budget deficit* (ex-post). However, if mechanisms that allocate inefficiently yield higher residual surplus (Guo and Conitzer 2008) it is not clear why one should use a mechanism that allocates efficiently.

Consequently, we drop the requirement that mechanisms allocate efficiently. Instead, we take an optimal mechanism design approach and consider mechanisms that are comparable to the ones considered before in that they are strategy-proof, deterministic, never run at a deficit and satisfy ex-post participation constraints. We analyze which mechanism maximizes residual surplus when an indivisible good is auctioned among two agents with independent private values that are distributed according to prior type distributions. We show that under an increasing hazard rate assumption on type distributions, the optimal mechanism will never waste any payments, thereby deviating distinctly from the efficient allocation (Theorem 1). In fact, our proof method reveals that all mechanisms that allocate efficiently are worse than the simple mechanism where the object is always given to one of the agents (Corollary 1), showing that our general mechanism design approach has clear advantages over the previous approach to search for the optimal Groves mechanism. We show that the optimal mechanism is either a "posted price" or an "option" mechanism: The object is assigned to one of the agents unless both agents agree to trade at a prespecified price (posted price mechanism) or unless the second agent uses his option to buy the object at a fixed price from the first agent (option mechanism). Therefore, the optimal mechanisms do not invoke money burning and are of a particularly simple form. In the bilateral trade setting, we establish optimality of posted price mechanisms without any restrictions on type distributions (Theorem 2). This provides an argument for the focus on budget-balanced mechanisms (see Myerson and Satterthwaite 1983, Hagerty and Rogerson 1987).

The requirement that a mechanism does not produce a budget deficit ex-post is considerably stronger than the requirement that this holds in expectation. However, in many situations it is reasonable that a budget breaker is infeasible and therefore ex-post constraints need to be obeyed. This includes situations where there is no insurance or where agents have restricted access to capital markets. Also, hidden information issues towards a third party cannot always be resolved, and autarkic mechanisms that can be implemented without explicit intervention by a third party might be preferable (e.g., when mechanisms are used to model bargaining situations). If all these considerations do not apply and mechanisms that create no deficit in expectation can be implemented, then one can achieve the first-best solution (see Section 5). Similarly, we show that one can potentially achieve the first-best if mechanisms are only required to be Bayesian incentive compatible. In contrast to these two constraints, which are the main driving forces behind our results, we argue that the participation constraint and the restriction to deterministic mechanisms are not essential to the spirit of our results.

Our work is part of a small literature that searches for mechanisms maximizing residual surplus when the first-best is not achievable. Miller (2012) studies a model of firms colluding in a Bertrand oligopoly. A mechanism used by a cartel to allocate market shares should maximize residual surplus. Miller shows that under general conditions it is never optimal to allocate market shares efficiently and gives numerical evidence that for some type distributions it is optimal to give up efficiency in order to obtain a balanced budget. However, other examples indicate that this observation does not hold for all distributions. Athey and Miller (2007) study residual surplus maximization in a repeated bilateral trade setting and obtain numerical results suggesting that for many type distributions the optimal mechanism is a posted price mechanism. Closely related to this chapter is independent work by Shao and Zhou (2012), who obtain the characterization of our Theorem 1 when restricting to symmetric distributions of types and allowing mechanisms to violate individual rationality.

Another related strand of the literature studies the expected residual surplus of Bayesian incentive compatible mechanisms when it is not possible to redistribute any payments among the agents (Hartline and Roughgarden 2008, Chakravarty and Kaplan 2013, Condorelli 2012). This implies that methods similar to those in Myerson (1981) can be applied. It is shown that for a large class of type distributions (those which exhibit an increasing hazard rate) it is optimal to always assign the object to the same agent. Maximization of residual surplus also plays a role in the analysis of optimal mechanisms used by bidding rings (McAfee and McMillan 1992). It is worth noting that the equivalence between Bayes-Nash and dominant strategy implementation (Gershkov, Goeree, Kushnir, Moldovanu and Shi 2013, Manelli and Vincent 2010) does not apply to our model.[1]

We present the basic model for the auction environment in Section 2 and characterize incentive compatible mechanisms in Section 3. The optimization problem is solved in Section 4, the role of the assumptions is discussed in Section 5. We study this mechanism design problem in the bilateral trade context in Section 6, and conclude in Section 7.

2. Model

An indivisible object is auctioned among two agents. Each agent $i = 1, 2$ has a valuation x_i for the object, which is his private information. Valuations are drawn independently from $X_i = [0, \bar{x}_i]$ according to distribution functions F_i with corresponding densities f_i, which we assume to be bounded. We denote by $X = X_1 \times X_2$ the product type space and by F the joint distribution on X. For notational convenience, when concentrating on agent i, we will write (x_i, x_{-i}) for $x = (x_1, x_2) \in X$.

If agent i is given a payment of p_i (usually negative), his utility is $x_i + p_i$ for winning the object, and p_i if the other agent gets the object.

Mechanisms

Due to the Revelation Principle we focus on truthfully implementable direct revelation mechanisms for selling the object.

Definition 1. *A* mechanism *M is a tuple (d, p), where $d : X \to \{0, 1\}^2$ and $p : X \to \mathbb{R}^2$ are measurable functions, such that $d_1(x) + d_2(x) = 1$.*[2]

[1]See Section 5 for more details.

[2]For a discussion of stochastic mechanisms, see Section 5. We follow Athey and Miller (2007) and Miller (2012) and assume that the good is always allocated. This is reasonable, for example, when considering how a cartel allocates market shares, or how the government sells licenses to firms. While there can be welfare gains from not allocating the good when one focuses on anonymous mechanisms (de Clippel, Naroditskiy, Polukarov, Greenwald and Jennings 2013), these gains seem to be minor in our model. Moreover, the assumption that the good is always allocated is without loss of generality in the trade setting (Section 6).

The interpretation is that $d_i(x) = 1$ if and only if agent i gets the object. If the agents report x, then agent i receives as payment the component $p_i(x)$ of $p(x)$.

Equilibrium Concept

We consider strategy-proof mechanisms where truthful reporting is a dominant strategy for both agents. Therefore, we define the following notion of incentive compatibility:

Definition 2. *A mechanism M is* incentive compatible (IC) *if for every agent i and for each $x_i \in X_i$, $r_i \in X_i$,*

$$d_i(x_i, r_{-i}) \cdot x_i + p_i(x_i, r_{-i}) \geq d_i(r_i, r_{-i}) \cdot x_i + p_i(r_i, r_{-i})$$

holds for each $r_{-i} \in X_{-i}$.

This definition is independent of the distribution of valuations, which reflects the robustness of strategy-proof mechanisms as compared to mechanisms that are Bayes-Nash incentive compatible. Although the set of mechanisms we consider does therefore not depend on F, the next section shows that the distributions determine which mechanism is optimal.

Objective and Further Constraints

We aim at finding the mechanism that maximizes the sum of agents' ex-ante (expected) residual surplus, that is, utility derived from the physical allocation minus aggregate payments. We impose the constraint that the mechanism has to be *ex-post no-deficit (ND)*, that is, for every type profile x, we require $p_1(x) + p_2(x) \leq 0$.[3] Also, the mechanism has to be *ex-post individually rational (IR)*, that is, for all type profiles x, we require $d_i(x)x_i + p_i(x) \geq 0, i = 1, 2$. Summarizing, we want to solve the following optimization problem:

$$\max_{M=(d,p)} \int_X \Big[d_1(x)x_1 + d_2(x)x_2 + p_1(x) + p_2(x) \Big] dF(x) \tag{1}$$
$$\text{s. t. } M \text{ satisfies IC, ND and IR.}$$

We say that a mechanism is optimal if it solves problem (1).

3. Characterization of Incentive Compatibility

The aim of this section is to give a characterization of incentive compatibility in order to simplify problem (1). The conditions characterizing incentive compatible mechanisms involve a monotonicity and an integrability condition. We first define monotonicity.

[3] Ex-post budget constraints are commonly imposed on mechanism design problems: see, for example, the literature on optimal redistribution (Guo and Conitzer 2010, Guo and Conitzer 2009, Moulin 2009) and bilateral trade (Hagerty and Rogerson 1987, Myerson and Satterthwaite 1983), or Chawla, Hartline, Rajan and Ravi (2006). The role of this assumption is discussed in Section 5.

Definition 3. *The allocation function d is* monotone *if d_i is non-decreasing in x_i for $i = 1, 2$.*

Now given a monotone allocation function d, define the following functions for $i = 1, 2$:

$$g_i(x_{-i}) := \inf\{x_i : d_i(x_i, x_{-i}) = 1\}.$$

If there is no x_i such that $d(x_i, x_{-i}) = 1$, then we set $g_i(x_{-i}) = \bar{x}_i$. Note that if d is monotone, these functions define d almost everywhere. The following lemma, which is a corollary of Myerson (1981), gives a characterization of incentive compatibility.

Lemma 1. *A mechanism $M = (d, p)$ is incentive compatible, if and only if the following two conditions are satisfied:*

1. *The allocation rule d is monotone.*
2. *For $i = 1, 2$ let $x_{-i} \in X_{-i}$ be given. Then for all $x_i \leq x'_i \in X_i$,*

$$p_i(x_i, x_{-i}) - p_i(x'_i, x_{-i}) = \begin{cases} g_i(x_{-i}) & \text{if } d_i(x_i, x_{-i}) = 0 \text{ and } d_i(x'_i, x_{-i}) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

The interpretation of condition (2) is that an agent who receives the object is punished by paying a higher amount compared to the case where he would not have gotten the object. The punishment has to make the agent's marginal type $g_i(x_{-i})$ indifferent between receiving and not receiving the object.

It follows from Lemma 1 that if a mechanism satisfies IC, payments have the following form:

$$p_i(x_i, x_{-i}) = q_i(x_{-i}) - g_i(x_{-i})d_i(x_i, x_{-i})$$

with some functions $q_i : X_{-i} \to \mathbb{R}$. This can be interpreted as a payoff-equivalence result: Payments are completely determined by the allocation as soon as one fixes the payment for some type x_i. Or, in other words, once the allocation is fixed, the only freedom that is left regarding the payment scheme, is to give the agent an additional payment that is independent of his type. These additional payments can serve as a possibility to redistribute certain amounts of payments to another agent, as e.g., in Cavallo (2006). Given an allocation rule d and a payment rule p, we say that the redistribution payment q implicitly defined by the above equality is associated with p.

The simplified formulation of problem (1) is the following:

$$\max_{M=(d,p)} \int_X \Big[d_1(x)[x_1 - g_1(x_2)] + d_2(x)[x_2 - g_2(x_1)] + q_1(x_2) + q_2(x_1)\Big] dF(x) \quad (3)$$
$$\text{s. t. } M \text{ satisfies IR and ND, } q \text{ is associated with } p \text{ and } d \text{ is monotone.}$$

We will write $U(M)$ for the above integral and from now on only consider mechanisms that are IC, IR and ND.

4. The Optimal Auction

In this section, we present the first main result of this chapter: if we impose an increasing hazard rate condition on the type distributions, then the optimal mechanism is always budget-balanced. Specifically, it turns out that the optimal mechanism takes one of two simple forms:

Either it is a posted price mechanism which by default allocates the object to one of the agents (agent 1, say) and changes the allocation if and only if both agents agree to trade at a prespecified price a, i.e., agent 1 reports a valuation below a fixed price a and agent 2 reports a valuation above a. If agent 2 is allocated the object, he makes a payment a to agent 1, otherwise no transfers accrue.

Or it is an option mechanism where the good is allocated by default to agent 1, but agent 2 has the option to buy the object at price a. Hence, if agent 2's valuation is above the strike price a, he buys the object and pays a to agent 1 (see also Shao and Zhou 2012).

Formally, these two mechanisms are defined as follows:

Definition 4. *A mechanism $M = (d, p)$ is a* posted price mechanism with default agent 1 and price a*, if*

$$\begin{array}{ll} d_2(x) = 1,\ p(x) = (a, -a) & \text{if } x_1 \leq a \text{ and } x_2 \geq a, \\ d_2(x) = 0,\ p(x) = (0, \quad 0) & \text{otherwise.} \end{array}$$

M is an option mechanism with default agent 1 and price a*, if*

$$\begin{array}{ll} d_2(x) = 1,\ p(x) = (a, -a) & \text{if } x_2 \geq a, \\ d_2(x) = 0,\ p(x) = (0, \quad 0) & \text{otherwise.} \end{array}$$

Similarly, one can define posted price and option mechanisms with default agent 2. If we do not specify the agent or price we just say that M is option or posted price.

Both classes of mechanisms are parameterized by the price a and it is easy to check that all these mechanisms are budget-balanced as well as incentive compatible and individually rational.

Our assumption on type distributions is the following:

Condition 1 (HR)**.** *The hazard rates of the type distributions are monotone. That is, the functions $h_i(x_i) = \frac{f_i(x_i)}{1-F_i(x_i)}$ are non-decreasing in $x_i \in [0, \bar{x}_i)$ for $i = 1, 2$.*

Theorem 1. *Given condition (HR), the optimal mechanism is either a posted price or an option mechanism.*

The proof can be sketched as follows: We first show the important auxiliary result that either an option mechanism or a posted price mechanism is optimal in $\mathcal{M}_0$, the class of mechanisms such that g_i is monotone and piecewise constant for each agent (that is, the line that separates the two allocation regions is a step function) (Lemma 2). We then argue that the welfare of a given mechanism can be approximated arbitrarily well by a mechanism in $\mathcal{M}_0$ (Lemma 3). The Theorem then follows by the following observation:

Suppose there is a mechanism $\bar{M}$ being strictly better than the best option or posted price mechanism, and denote the welfare difference by ε. It follows from Lemma 3 that there is a mechanism in the class $\mathcal{M}_0$ whose welfare is within $\frac{\varepsilon}{2}$ of $U(\bar{M})$, thus being better than the best option or posted price mechanism. But this contradicts Lemma 2, hence there cannot be a mechanism being better than the best option or posted price mechanism.

While the approximation part of the proof can be found in the appendix, we state and prove Lemma 2, which contains the essence of why Theorem 1 holds.

Lemma 2. *Assume condition (HR) and let $M = (d, p)$ be any mechanism in $\mathcal{M}_0$. Then there exists a mechanism M' that is posted price or option such that $U(M') \geq U(M)$.*

Proof. The proof consists of three steps, where we constructively manipulate M in order to end up with the desired mechanism M'. We denote the jump points of $g_2(x_1)$ and $g_1(x_2)$ by α_j and β_j, respectively (see Figure 1).

Step 1: This step shows how to determine the maximal possible redistribution payments q_i. To this end, we note that, without loss of generality, we can assume that for the first segment of g_1 we have $g_1(x_2) = 0$ since otherwise we could switch the roles of the agents.

We now claim that $q_2(x_1) = 0$, $\forall x_1 \in X_1$; that is, no money is redistributed to agent 2. To see this, pick arbitrary x_1 and observe that $g_1(0) = 0$ and $d_2(x_1, 0) = 0$; therefore $g_1(0)d_1(x_1, 0) = g_2(x_1)d_2(x_1, 0) = 0$. From (ND) it follows that $q_1(0)+q_2(x_1) = p_1(x_1, 0) + p_2(x_1, 0) \leq 0$. Also, (IR) for agent 2 at $(x_1, 0)$ implies $q_2(x_1) \geq 0$, and (IR) for agent 1 at $(0, 0)$ implies $q_1(0) \geq 0$, and therefore $q_2(x_1) = 0$.

Next, we can assume that

$$q_1(x_2) = \min_{x_1} \left\{ g_1(x_2)d_1(x_1, x_2) + g_2(x_1)d_2(x_1, x_2) \right\} \tag{4}$$

always holds, since by (ND) this relation always holds with $\leq$ and changing it to equality does not reduce $U(M)$. In this way, the complete payment-scheme is determined through the allocation rule d. Note that setting the function q this way implies that (ND) and (IR) are always satisfied.

Step 2: In this step we argue that changing the allocation to the one shown in Figure 1b does not increase money burning, but increases allocative efficiency and hence aggregate welfare.

Define the set $B = \{x \mid x_1 \leq \beta_1 \leq x_2, d_2(x) = 0\}$ and consider the sets B_1, B_2 and C as shown in Figure 1a. We change the allocation rule and allocate the object to agent 2 for types in B. Since $x_2 \geq x_1$ for $x \in B$, this improves the physical allocation and we can concentrate on payments. Note that q_1, as defined in (4), increases to the same extent as g_1, hence any additional payments in the set B_2 can be redistributed. Also, transfers are weakly increased for types in B_1 and C. As the change in allocation has no effect outside these sets, the claim follows.

Step 3: This step studies the effects of shifting steps in the set R, shown as the shaded area in Figure 1b, while fixing redistribution payments. Our condition on the hazard rate ensures that each step should optimally be moved to either the lowest or

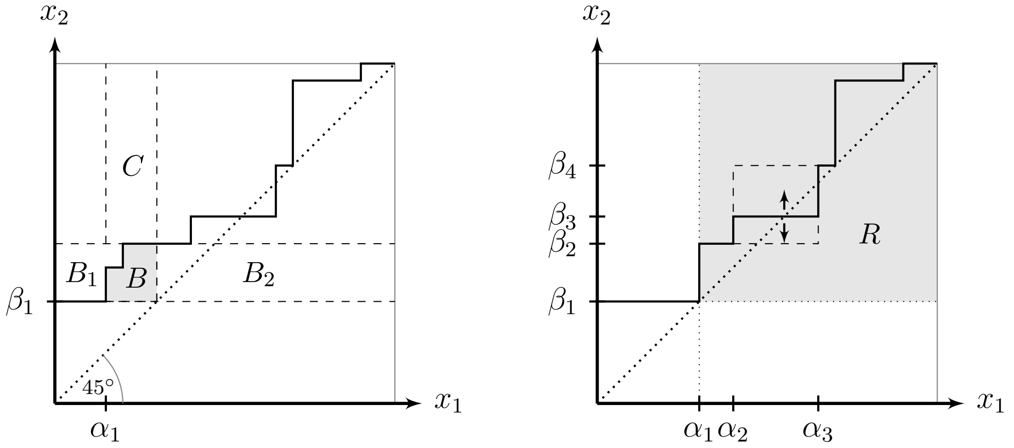

(a) Illustration of the sets used in Step 2 of the proof.

(b) Situation after the object is assigned to agent 2 in the set B. The dashed lines indicate the interval in which the jump point in Step 3 is varied.

Figure 1: Illustration of the proof of Lemma 2.

the highest possible position. Hence, proceeding iteratively, we obtain either an option mechanism or a posted price mechanism. This will complete the proof.

Changing the allocation in R does not change q_1 as defined in (4) and we ignore the functions q_i from now on.

The following is a procedure to remove one step contained in R without decreasing $U(M)$. We do this exemplarily with the jump point at β_3 (see Figure 1b). We vary β_3 on the interval $[\beta_2, \beta_4]$ and show that welfare is quasi-convex in β_3. This implies that setting $\beta_3^* = \beta_2$ or β_4 increases $U(M)$. The part of $U(M)$ that depends on β_3 is the following:

$$\int_{\alpha_2}^{\alpha_3} \left[\int_{\beta_2}^{\beta_3} (x_1 - \alpha_2) dF_2(x_2) + \int_{\beta_3}^{\bar{x}_2} (x_2 - \beta_3) dF_2(x_2) \right] dF_1(x_1)$$
$$- \int_{\alpha_3}^{\bar{x}_1} \left[\int_{\beta_2}^{\beta_3} \alpha_2 dF_1(x_1) + \int_{\beta_3}^{\beta_4} \alpha_3 dF_2(x_2) \right] dF_1(x_1)$$

Differentiating with respect to β_3 using Leibniz' rule yields

$$\int_{\alpha_2}^{\alpha_3} \Big[f_2(\beta_3)(x_1 - \alpha_2) - \big[1 - F_2(\beta_3)\big] \Big] dF_1(x_1) + \int_{\alpha_3}^{\bar{x}_1} f_2(\beta_3)[\alpha_3 - \alpha_2] dF_1(x_1).$$

Writing constants C_1, C_2 and C_3 for the terms that do not depend on β_3, we get

$$C_1 f_2(\beta_3) - C_2 \big[1 - F_2(\beta_3)\big] + C_3 f_2(\beta_3).$$

Assuming $C_2[1 - F_2(\beta_3)] > 0$ (if either $C_2 = 0$ or $1 - F_2(\beta_3) = 0$, we set $\beta_3^* = \beta_4$ without reducing U), we can divide by $C_2[1 - F_2(\beta_3)]$ and get that the derivative is non-negative if and only if

$$C \cdot h_2(\beta_3) - 1 \geq 0,$$

where $C = (C_1 + C_3)/C_2 > 0$. Because $h_2(\beta_3)$ is non-decreasing by condition (HR), quasi-convexity follows and $U(M)$ is increased by either setting $\beta_3^* = \beta_2$ or $\beta_3^* = \beta_4$. In either case, we have decreased the number of steps by one and the procedure ends.

Iteratively applying this procedure establishes the lemma. □

A consequence of the theorem is that, given the increasing hazard rates of the agents' type distributions, finding the best mechanism reduces to finding the best posted price and option mechanisms and comparing these two. For example, if the agents have the same distribution function, all option and posted price mechanisms with the same strike price yield the same welfare and therefore the best mechanism is characterized by the strike price a^* satisfying

$$a^* = \mathbb{E}[x_1] = \mathbb{E}[x_2].$$

Our intermediate results (see the proof of Lemma 2) also allow for a refined judgment of the welfare implied by the efficient allocation. Miller (2012) showed, under very general conditions, that the efficient allocation rule is never part of the optimal mechanism. We can strengthen this statement in our context by providing a mechanism that improves upon all efficient mechanisms. Surprisingly, this improvement can be achieved using an extremely simple mechanism:

Corollary 1. *Given condition (HR), every mechanism that allocates efficiently is dominated by a mechanism that always allocates the good to one of the agents.*

More precisely, a mechanism that is better than every efficiently allocating mechanism can be found simply by comparing the agents' type distributions, giving the good to the agent with the higher expected valuation and completely ignoring any reported types.

While optimal mechanisms for distributions obeying condition (HR) are very simple, the following example shows that if the condition is not satisfied the optimal mechanism need not be of the form stated in Theorem 1. The example also illustrates the role of (HR) in establishing the result.

Example 1. *Let the distribution function of two symmetric agents be given as*

$$f(x_i) = \begin{cases} 0.9 & \text{if } x_i \leq 0.5 \\ 0.1 & \text{otherwise.} \end{cases}$$

Due to the downwards jump at 0.5*,* f *does not satisfy condition (HR). The optimal posted price mechanism (which is as good as the optimal option mechanism) has a strike price of* $a^* = 0.275$*, attaining a social welfare of* 0.0718*. However, the following mechanism* M *attains a higher social welfare of* 0.0741*: Set*

$$d_2(x) = 1 \quad \Leftrightarrow \quad (x_2 \geq a^* \text{ and } x_1 \leq a^*) \text{ or } (x_2 \geq 0.5 \text{ and } x_1 \leq 0.5),$$

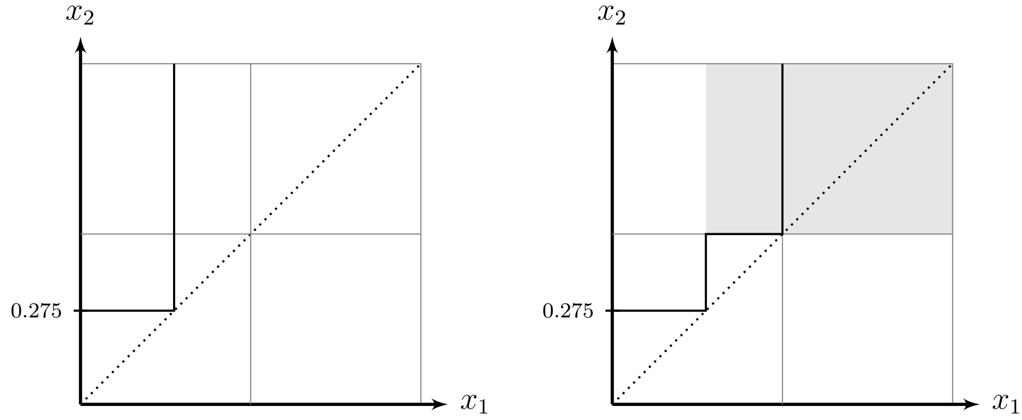

(a) Optimal posted price mechanism.

(b) Mechanism M.

Figure 2: Mechanisms presented in Example 1

and set $q_2(x_1) \equiv 0$, *as well as*

$$q_1(x_2) = \begin{cases} 0 & \text{if } x_2 \leq a^* \\ a^* & \text{otherwise.} \end{cases}$$

This mechanism and the best option mechanism are depicted in Figure 2. One can see that the allocation of mechanism M *is more efficient. Because the induced higher payments cannot be redistributed, payments of* $(0.5 - 0.275) = 0.225$ *are lost for type profiles in the shaded area in Figure 2b. But still, since type profiles* x *with* $x_1, x_2 \geq 0.5$ *appear so rarely (with density* 0.01*), this does not counter the positive effect due to the better allocation. In this sense, an increasing hazard rate ensures that lost payments can never be weighed out by an improved efficiency of the allocation.*

5. Robustness

In this section, we in turn relax the no-deficit, incentive compatibility and participation constraints as well as the restriction to deterministic mechanisms, and analyze how sensitive our characterization in Theorem 1 is to these relaxations.

Ex-ante Budget Constraints

While ex-post budget constraints are imposed commonly in the literature and seem appropriate for many settings, they would effectively be turned into ex-ante constraints if insurance against budget deficits was available.[4] Relaxing the no-deficit constraint to an ex-ante constraint, thus requiring the mechanism to run at no deficit on average, simplifies the problem and allows the planner to implement the first-best. This can be achieved by running the VCG mechanism. This mechanism is ex-post individually

[4]Note also, that the exact form of the budget constraints can be irrelevant when considering Bayesian incentive compatible mechanisms (Esö and Futo 1999).

rational and creates no deficit ex-post. By redistributing the expected surplus in an arbitrary fixed way to the agents, the mechanism becomes ex-ante budget-balanced and therefore achieves the first-best.

Bayesian Incentive Compatible Mechanisms

If stronger assumptions can be made on the information structure (namely, if the agents' beliefs equal a common prior that is known to the designer), we can relax the constraints on the mechanisms to Bayesian incentive compatibility and interim individual rationality. This allows the implementation of mechanisms that achieve higher expected welfare. Notably, if the distribution of types is symmetric across agents, then the expected externality mechanism (d'Aspremont and Gerard-Varet 1979, Arrow 1979) achieves the first-best.[5] To see this, observe that this mechanism allocates efficiently, has a balanced budget, and has payments given by

$$t_i(\theta) = \int_{\theta_i}^{\overline{\theta}_{-i}} \theta_{-i} dF_{-i}(\theta_{-i}) - \int_{\theta_{-i}}^{\overline{\theta}_i} \theta_i dF_i(\theta_i).$$

Therefore, an agent reporting a type of 0 receives a weakly positive transfer and hence a weakly positive utility.

This implies that the equivalence of Bayesian and dominant strategy incentive compatible mechanisms established by Gershkov, Goeree, Kushnir, Moldovanu and Shi (2013) does not apply. They show that in a large class of mechanism design problems, for any Bayesian incentive compatible and interim individually rational mechanism, there exists an equivalent dominant strategy incentive compatible mechanism that is ex-post individually rational. However, this equivalence is established in the absence of budget constraints, and the above arguments imply that it cannot be extended to our setting.

Participation Constraints

While our general characterization of the optimal mechanism does not hold with relaxed participation constraints, these constraints are not the main driving force behind our results and the inefficiency of the optimal allocation. Indeed, our characterization can be obtained without participation constraints if one restricts attention to settings where agents are symmetric ex-ante (Shao and Zhou 2012).

Stochastic Mechanisms

In the previous section we restricted attention to deterministic mechanisms in order to be able to analytically characterize the optimal mechanism. Deterministic mechanisms have additional benefits: They are simpler to implement, and more plausible in some settings (e.g., when modeling bargaining between agents).

[5] More generally, a mechanism in the spirit of Myerson and Satterthwaite (1983) is optimal, that allocates the object to the agent with the highest weighted virtual valuation.

Distribution	Average loss	Maximum loss	Instances without loss
Random	0.018 %	0.874 %	92.500 %
IHR	0.003 %	0.420 %	97.850 %
Weibull	0.000 %	0.000 %	100.000 %
All	0.007 %	0.874 %	96.785 %

Table 1: Simulation results comparing the welfare loss due to the restriction to deterministic mechanisms.

While there are instances where the focus on deterministic mechanisms is not without loss, numerical simulations suggest that the induced loss in welfare is small. We generated $n = 2\,000$ random instances for three classes of distributions of types: Random distributions, random distributions with an increasing hazard rate, and distributions from the Weibull class with different shape and scale parameters such that the distribution has an increasing hazard rate. We then computed the optimal deterministic and stochastic mechanism for every instance. The results are summarized in Table 1 which shows, for each distribution class, the average and maximum welfare loss of the optimal deterministic mechanism, as a percentage of the welfare of the best stochastic mechanism. The fourth column shows the percentage of instances where there is no loss due to the restriction to deterministic mechanisms. As can be seen, instances where the deterministic constraint is binding appear only rarely. Further, even if this is the case, the percentage loss in expected welfare is very small.

6. Bilateral Trade

Myerson and Satterthwaite (1983) showed that one cannot implement the efficient allocation in the bilateral trade setting in an ex-post budget-balanced and interim individually rational way, and characterized the optimal mechanism satisfying these constraints. In the same environment, Hagerty and Rogerson (1987) study the set of dominant-strategy implementable mechanisms that are ex-post budget-balanced and individually rational, showing that essentially only posted price mechanisms fulfill these conditions. However, a priori it is not clear why one should restrict the search for the optimal mechanism to mechanisms with a balanced budget. After all, it is conceivable that deviating from a balanced budget could improve incentives and therefore lead to higher welfare. In fact, Schwartz and Wen (2012) show by example that relaxing budget-balancedness to a no-deficit constraint can improve upon posted price mechanisms. The result in this section shows that this holds only for stochastic mechanisms; when looking at deterministic mechanisms, the restriction to budget-balanced mechanisms does not reduce aggregate welfare.

Let the model and notation be as in Section 2, but assume now that agent 1 (called the "seller" from now on and indexed by S) is the owner of the good before participating in the mechanism (whereas agent 2 is called the "buyer" and indexed by B). By a buyer posted price mechanism (B-PP) we denote a posted price mechanism in which the buyer gets the object if and only if he announces a type high enough, and the seller a type

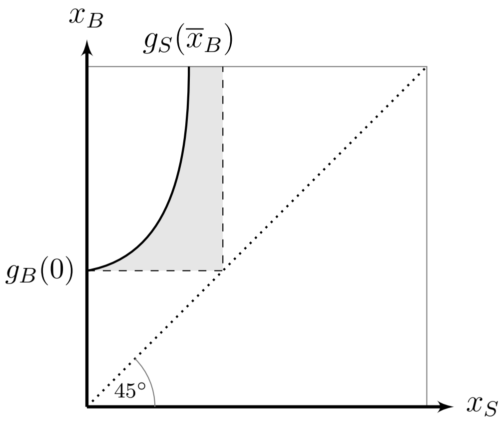

Figure 3: Illustration of the proof of Theorem 2. The shaded area indicates the type profiles where the initial mechanism differs from the posted price mechanism with strike price $g_B(0)$ (dashed line).

that is low enough. Again, we are looking for a mechanism that maximizes the sum of the expected utilities of the agents, taking monetary transfers into account. The fact that in the bilateral trade setting the seller initially owns the good requires a stronger condition for a mechanism to be individually rational: now the outside option for a seller is to not participate in the mechanism and to keep the object. Hence, for a mechanism to be individually rational,

$$d_S(x)x_S + p_S(x) \geq x_S \qquad \text{and} \qquad d_B(x)x_B + p_B(x) \geq 0 \tag{IR'}$$

must hold for all $x \in X$.

Thus, a mechanism is optimal if it solves

$$\max_{M=(d,p)} \int_X \Big[d_S(x)x_S + d_B(x)x_B + p_S(x) + p_B(x)\Big] dF(x) \tag{5}$$
$$\text{s.t. } M \text{ satisfies IC, ND and IR'}.$$

Theorem 2. *There is a B-PP mechanism that solves problem* (5).

Proof. We first show that (IR') implies that the seller keeps the object whenever his valuation is higher. Assume to the contrary that trade takes place at $x_S > x_B$. Then (IR') for the seller implies that the seller receives at least x_S and (IR') for the buyer implies that he pays at most x_B, violating (ND).

Recall that $g_B(0)$ denotes the smallest buyer type such that trade takes place when $x_S = 0$, and $g_S(\overline{x}_B)$ denotes the highest seller type such that trade takes place when $x_B = \overline{x}_B$. We claim that $g_S(\overline{x}_B) \leq g_B(0)$. Constraints (IC) and (IR') for the seller imply that he receives at least a payment of $g_S(\overline{x}_B)$ whenever the buyer reports $\overline{x}_B$ and trade takes place, in particular at $(0, \overline{x}_B)$. Similarly, (IC) and (IR') for the buyer imply that he pays at most $g_B(0)$ whenever the seller reports 0, in particular at $(0, \overline{x}_B)$. Therefore, $g_S(\overline{x}_B) > g_B(0)$ would violate (ND) at $(0, \overline{x}_B)$.

Finally, we claim that the B-PP mechanism with strike price $g_B(0)$ weakly dominates the given mechanism. To see this, note that $p_S(x) + p_B(x) \leq 0$ by (ND) and a B-PP mechanism is budget-balanced. Hence, the posted price mechanism dominates the old mechanism with respect to payments. Since the allocation only differs for x such that $x_B \geq g_B(0) \geq x_S$ and the posted price allocation rule prescribes $d_B(x) = 1$ for such x (see also Figure 3), the posted price mechanism also dominates the old mechanism with respect to the allocation rule. □

In contrast to Theorem 1, this result shows that a posted price mechanism is optimal for any type distribution. The difference is due to the stronger individual rationality constraint. While any allocation rule is compatible with (IR), the stronger constraint (IR') in the trade setting restricts the set of allocation rules that can be implemented without a budget deficit. Within this smaller class of feasible allocation rules, for any distribution of types a posted price mechanism is optimal.

The stronger individual rationality constraint also implies that mechanisms which do not allocate the object are infeasible. This is because if the buyer does not get the object, no money can be collected to compensate the seller for losing the object. Therefore, assuming that the object is always allocated is without loss of generality in this setting.

7. Discussion

We have studied the trade-off between efficiency and budget-balancedness in an independent private values auction model. We incorporated this into the model by letting the social welfare objective function include all payments, that is, by maximizing residual surplus.[6] We showed that, if one focuses on robust implementation in dominant strategies, an increasing hazard rate condition on agents' type distributions guarantees a resolution of the trade-off completely in favor of a balanced budget. In addition, budget-balanced mechanisms have a very simple form and can easily be implemented as posted price or option mechanisms. Further, we showed without any assumption on the prior distribution of types that a posted price mechanism is optimal in the bilateral trade setting. Our results imply that our approach of optimal mechanism design yields higher welfare than approaches concentrating on the efficient allocation.

In the section on robustness we have seen that while the restriction to deterministic and ex-post individually rational mechanisms is not crucial for our main result, it is primarily driven by the focus on strategy-proof mechanisms that satisfy the ex-post no-deficit constraint: Without these constraints, in many settings the first-best can be achieved. This shows that these two restrictions are relatively costly in terms of welfare.

An interesting open question is how the result generalizes to a model including more than two agents. We strongly believe that the optimal mechanism will still be budget-balanced. An important argument for this is that, as the number of agents gets large, the efficient allocation can be approximated in a budget-balanced way: in the spirit of McAfee (1992), allocate efficiently while ignoring one agent who then receives all payments from the other agents. This can be implemented by tentatively giving the

[6] For other ways to analyze the frontier that describes possible ways to resolve this trade-off, see, for example, Diakonikolas, Papadimitriou, Pierrakos and Singer (2012) or Tatur (2005).

object to one of the agents and then simulating a second price auction with reserve price where this agent sells the object to the remaining agents.

Appendix

The following lemma enables us to approximate any mechanism with mechanisms from the class $\mathcal{M}_0$.

Lemma 3. *For every mechanism $M = (d, p)$ and for every $\varepsilon > 0$ there exists a mechanism $\tilde{M} = (\tilde{d}, \tilde{p})$ in $\mathcal{M}_0$ such that $U(M) - U(\tilde{M}) < \varepsilon$.*

Proof. Let the mechanism $M = (d, p)$ and $\varepsilon > 0$ be given and let $g_1(x_2)$ and $g_2(x_1)$ be defined as above. Define $D_i := \{x \in X : d_i(x) = 1\}$ as the set of type profiles where agent i gets the object and define $\tilde{D}_i$ similarly. Since g_2 is a monotone function it can be approximated uniformly by a monotone step function $\tilde{g}_2$. Denote the associated allocation rule by $\tilde{d}$. By choosing the step width small enough the approximation can be done such that for given $\delta > 0$,

$$||g_1 - \tilde{g}_1||_\infty < \delta \quad \text{and} \quad ||g_2 - \tilde{g}_2||_\infty < \delta$$

holds. The approximation can be chosen such that $g_i(x_{-i}) = \bar{x}_i$ implies $\tilde{g}_i(x_{-i}) = \bar{x}_i$ and $\tilde{g}$ can be chosen such that $\tilde{g}_2 \leq g_2$, implying that $\tilde{D}_1 \subset D_1$.

Without loss of generality, we can assume that $q_2(x_1) \equiv 0$ (see Step 1 in the proof of Lemma 2). By construction of $\tilde{g}_2$ and since M satisfies (ND), we can define functions $\tilde{q}_i(x_{-i})$ such that $\tilde{q}_2(x_1) \equiv 0$, $0 \leq \tilde{q}_1(x_2) \leq \inf_{x_1}\{\tilde{g}_1(x_2)\tilde{d}_1(x_1, x_2) + \tilde{g}_2(x_1)\tilde{d}_2(x_1, x_2)\}$ $\forall x_2 \in X_2$ and $||\tilde{q}_1 - q_1||_\infty < \delta$. We then have:

$$\begin{aligned}
U(d,p) - U(\tilde{d},\tilde{p}) &\leq \int_X q_1(x_2) - \tilde{q}_1(x_2)\ dF(x) \\
&\quad + \int_{D_1} x_1 - g_1(x_2)\ dF(x) - \int_{\tilde{D}_1} x_1 - \tilde{g}_1(x_2)\ dF(x) \\
&\quad + \int_{D_2} x_2 - g_2(x_1)\ dF(x) - \int_{\tilde{D}_2} x_2 - \tilde{g}_2(x_1)\ dF(x) \\
&\leq \delta + \int_{D_1 \backslash \tilde{D}_1} x_1 - g_1(x_2)\ dF(x) + \int_{\tilde{D}_1} \delta\ dF(x) \\
&\quad + \int_{\tilde{D}_2 \backslash D_2} x_2 - g_2(x_1)\ dF(x) + \int_{D_2} \delta\ dF(x) \\
&\leq 3\delta + B_1\overline{x}_1\delta + B_2\overline{x}_2\delta,
\end{aligned}$$

where B_i is an upper bound for $f_i(x_i)$. Hence, by choosing $\delta < \frac{\varepsilon}{3+B_1\overline{x}_1+B_2\overline{x}_2}$, it follows that $U(d,p) - U(\tilde{d},\tilde{p}) < \varepsilon$. □

We combine the approximation lemma with Lemma 2 in order to prove the theorem.

Proof of Theorem 1. Without loss of generality, we restrict ourselves to posted price mechanisms for agent 2. We first establish that U maps the set of all posted price mechanisms to a compact subset of $\mathbb{R}$. Let $\bar{a} = \min\{\bar{x}_1, \bar{x}_2\}$ and let $a \in [0, \bar{a}]$ be some price for a posted price mechanism M_a. Then $U(M_a)$ can be written as

$$U(M_a) = \int_0^a \int_a^{\bar{x}_2} x_2 dF(x) + \int_0^{\bar{x}_1} \int_0^a x_1 dF(x) + \int_a^{\bar{x}_1} \int_a^{\bar{x}_2} x_1 dF(x).$$

Due to the continuity of F, this function is continuous with respect to a. Since $[0, \bar{a}]$ is compact, so is $\{U(M_a) \mid a \in [0, \bar{a}]\}$ and therefore there exists an a^* such that $U(M_{a^*})$ is maximal among all posted prices.

Next, assume that the theorem is false, i.e., there exists a mechanism M and $\varepsilon > 0$ such that $U(M) > U(M_{a^*}) + \varepsilon$. Then apply Lemma 3 to M and ε to get a mechanism $\tilde{M} \in \mathcal{M}_0$ with $U(\tilde{M}) > U(M_{a^*})$. This contradicts Lemma 2, establishing the theorem. □

Chapter 4

Substitutes and Complements in Trading Networks

We generalize the full substitutes condition used in the trading network model of Hatfield et al. (2013) to a condition that we call full substitutes and complements (see Sun and Yang 2006). If all agents' preferences satisfy full substitutes and complements, competitive equilibria can be shown to exist and all desirable results about competitive equilibria carry over to the model with more diverse preferences: The welfare theorems hold and, under the full substitutes and complements condition, competitive equilibrium outcomes are precisely those that are stable.

1. Introduction

This chapter builds on the trading network model introduced by Hatfield et al. (2013). Their model is itself based on a hierarchy of models that analyze the concepts of competitive equilibrium, core and stability in markets with indivisible items: In the assignment model (Shapley and Shubik 1971, Gale 1960, Koopmans and Beckmann 1957), there is a set of sellers who each have an item for sale and a set of potential buyers with a willingness to pay for each item. The efficient allocation of items can be supported by Walrasian equilibrium prices and the outcomes induced by these prices correspond exactly to the stable outcomes and the core of the assignment game.

When buyers express their preferences over whole bundles of items, the *gross substitutes* property ensures that competitive equilibrium prices supporting the efficient allocation of items continue to exist (Kelso and Crawford 1982). Gul and Stacchetti (1999) show that these prices can be chosen such that they are anonymous in the sense that different buyers would pay the same price for a given item. The gross substitutes condition says that buyers view the items as substitutes, requiring that if the price for some item rises a buyer will not decrease his demand for all other items. Ostrovski (2008) recognized that the bilateral structure inherent in the previous auction and matching models can be extended to supply chains: Stable outcomes and equilibria still exist if agents view trades on the same side of the market as substitutes for each other but upstream trades as complementing downstream trades. Hatfield et al. (2013) generalize the trading relationships to arbitrary network structures where every agent may engage in various sell and/or buy relationships with other agents.

Sun and Yang (2006) generalize the model of Kelso and Crawford (1982) in a different way. They show that there is a condition on preferences that is more general than gross substitutes and still guarantees the existence of Walrasian equilibria: The items to be sold can be divided into two classes (tables and chairs) and all buyers view items within a class as substitutes for each other and items across classes as complements (*gross substitutes and complements*). Since the partition into classes is the same for each buyer, this result could escape known necessity results for the existence of Walrasian equilibria (Milgrom 2000, Gul and Stacchetti 1999).

This chapter unifies both generalizations of the Kelso and Crawford (1982) economy into one model while escaping the necessity result for the existence of Walrasian equilibria given in Hatfield et al. (2013). We show that there is an analogue version of the gross substitutes and complements property tailored to the trading network economy (*full substitutes and complements*), such that a Walrasian equilibrium is guaranteed to exist.

The model considered in this chapter is the same as in Hatfield et al. (2013): There is a set of agents who each may engage in a set of possible trading relationships in which they are either the seller or the buyer. The agents have cardinal preferences over subsets of the possible trades. A price can be associated with each trade and this induces a quasi-linear utility function for each agent. The interpretation of the full substitutes condition introduced by Hatfield et al. (2013) is that agents view the items underlying all possible trades as substitutes for each other. This means that if the price for some item that an agent is buying increases, she will increase her demand for items in other trades that she is currently involved in: She will buy more and sell less of the other items. Thus, the fact that an agent takes the role of a seller in one trade and that of a buyer in another trade leads to an apparent complementarity between trades on different sides of the market. However, the items underlying the trades are substitutes for the agent. The full substitutes and complements condition introduced in this chapter generalizes this in the same way the gross substitutes and complements condition (Sun and Yang 2006) generalizes gross substitutes: The items underlying all possible trades are partitioned into two classes, so that there are, e. g., "software" and "hardware" products. Buyers view different software products as being substitutes to each other but complemented by hardware. This means that if the price for a piece of hardware that an agent buys rises, he will weakly increase his demand for other hardware and weakly decrease his demand for software. In terms of trading relationships, this means that the agent will buy more and sell less hardware, and buy less and sell more software. Similar as above, this leads to apparent substitutability between hardware trades on one side of the market and software trades on the other side of the market although the underlying items are complements to each other.

We are able to show that when all agents' preferences satisfy the full substitutes and complements condition there always exists a Walrasian equilibrium. The proof for this result has the same structure as the proof in Hatfield et al. (2013): For a given network, a reduced bilateral economy is constructed in which trades are treated as the items that are for sale. The core of the proof then consists of showing that preferences in the reduced economy satisfy the gross substitutes and complements condition.

Applying the results in Sun and Yang (2006) implies that in the reduced economy a Walrasian equilibrium exists which can be lifted back to the original economy. The way we generalize preferences circumvents the necessity result for the existence of competitive equilibria since its proof relies on constructing preferences for agents who do not all view the same items as complementing each other.

We also show that the results in Hatfield et al. (2013) regarding the connection between competitive equilibrium and stability carries over to our assumption on preferences. While competitive equilibrium outcomes are always stable, the assumption of full substitutes and complements preferences ensures that the allocation of every stable outcome can be supported in Walrasian equilibrium. Furthermore, we show that if an agent is indifferent about the identity of his trading partners (i.e., he wants to sell an item but does not care about who the buyer is) then there exists a competitive equilibrium in which all these trades occur at the same price. This implies that all bilateral auction models (Sun and Yang 2006, Gul and Stacchetti 1999) are strictly contained in our setting because the result implies that in competitive equilibrium, every item that is for sale can be assigned exactly one price.

Recently, efforts have been successful towards a full characterization of the types of preferences such that a competitive equilibrium exists (Baldwin and Klemperer 2013). While this result can alternatively be used to derive our main result, the characterization is a geometric condition on the valuation function and not easy to interpret economically.

The chapter is structured as follows: In Section 2 we present the model and formally introduce the full substitutes and complements condition. Section 3 presents the existence result, describes the reduction to the bilateral economy and also explains how the result relates to previous reductions to the Kelso and Crawford (1982) model. In Section 4 we deal with the construction of anonymous competitive equilibrium prices and analyze the relationship between the concepts of competitive equilibrium and stability. All proofs related to Section 4 are relegated to the appendix. Finally, Section 5 presents a discussion of the results.

2. Environment

There is a set of agents I and a set of possible trades Ω which are exogenously given. A *trade* $\omega \in \Omega$ is an ordered pair of agents $\omega = (s_\omega, b_\omega)$, where s_ω and b_ω are the *seller* and the *buyer* of the trade ω, respectively. Thus, the agents and the set of trades can be thought of as a directed multigraph, where I is the set of vertices and Ω is the set of edges.

The trades Ω are partitioned into two classes, $\Omega = \Omega^1 \sqcup \Omega^2$, where we think of Ω^1 as the possible software trades and of Ω^2 as the set of possible hardware trades. A *price vector* $p \in \mathbb{R}^\Omega$ specifies a price for every trade $\omega \in \Omega$, and we write $p_\omega = p(\omega)$. An *arrangement* (Ψ, p) is a subset of trades $\Psi \subseteq \Omega$ together with a price vector p.

Some notation that will be used quite often: Let $A \subseteq \Omega$. Then we write $A^1 = A \cap \Omega^1$ for the software trades in A, $A_{\to i} = \{(s,b) \in A \mid b = i\}$ for the trades in A that agent i buys, $A_{i\to} = \{(s,b) \in A \mid s = i\}$ for the trades in A that agent i sells, $A_i = A_{\to i} \cup A_{i\to}$ for all trades agent i is involved in, and $A^1_{i\to} = A_{i\to} \cap \Omega^1$ for the software trades in A

that agent i sells (similarly, define $A^1_{\to i}$ and A^1_i).

Preferences

Every agent i has a *valuation function* $u_i : 2^{\Omega_i} \to \mathbb{R}$ over sets of trades $A_i \subseteq \Omega_i$.[1] Together with a price vector p, the valuation function induces the quasilinear *utility function* U_i over sets of trades A and prices p, defined as

$$U_i(A, p) := u_i(A_i) + \sum_{\omega \in A_{i\to}} p_\omega - \sum_{\omega \in A_{\to i}} p_\omega.$$

For convenience, we will sometimes write

$$p(A_{i\to}) = \sum_{\omega \in A_{i\to}} p_\omega \quad \text{and} \quad p(A_{\to i}) = \sum_{\omega \in A_{\to i}} p_\omega,$$

as well as $p(A) = p(A_{i\to}) - p(A_{\to i})$. The utility function U_i allows us to define agent i's *demand correspondence*, specifying which sets of trades maximize the agent's utility, given prices p:

$$D_i(p) := \arg\max_{A \subseteq \Omega_i} U_i(A, p)$$

For our definition of the full substitutes and complements condition we will use an analogue of the *indicator language full substitutes (IFS)* property of Hatfield et al. (2013), since it is the easiest to understand and interpret.[2] To this end, define for each agent i the indicator function $e(A) \in \{-1, 0, 1\}^{\Omega_i}$ for a set of trades $A \subseteq \Omega_i$ as

$$e_\omega(A) = \begin{cases} -1 & \text{if } \omega \in A_{i\to} \\ 0 & \text{if } \omega \notin A \\ 1 & \text{if } \omega \in A_{\to i} \end{cases} \qquad \omega \in \Omega_i.$$

The indicator function notation can be interpreted such that the agent demands a positive amount of an item if it is contained in a buy trade, and that he demands a negative amount of the item if it is contained in a sell trade.

Definition 1. *A valuation function u_i satisfies the* full substitutes and complements (FSC) *property, if the following conditions hold:*

1. *Take two price vectors $p \le q$, where $p_\omega = q_\omega$ for $\omega \in \Omega^2$. Then for all $A \in D_i(p)$ there exists $B \in D_i(q)$ such that*

$$\begin{aligned} e_\omega(A) &\le e_\omega(B) \quad \text{for all } \omega \in \Omega^1_i \text{ with } p_\omega = q_\omega, \\ e_\omega(A) &\ge e_\omega(B) \quad \text{for all } \omega \in \Omega^2_i. \end{aligned} \tag{1}$$

[1] For simplicity, we do not allow the valuation functions to take on the value minus infinity, which would model technological constraints on the set of trades an agent may engage in. However, doing this is with no loss of generality, since, similarly as in Hatfield et al. (2013), one can transform an unbounded valuation function to one of those we consider.

[2] In order to keep this chapter short, we will not delve into showing equivalence to the various analogues of *choice language full substitutability* and *demand language full substitutability*. All proofs presented in this chapter are adapted to use only the definition given below.

2. The above condition also holds if the sets Ω^1 and Ω^2 are exchanged.

The condition requires that, if prices for software trades go up, then the agent increases his demand for other software products whose price remained constant (because they are substitutes), and the agent decreases his demand for hardware products (because they are complements to software products that got more expensive). Note that for the case where $\Omega = \Omega^1$ and $\Omega^2 = \emptyset$, FSC reduces to the full substitutes condition.

3. Existence of Competitive Equilibria

In this section we prove that in an economy where each agent's preference satisfies the FSC condition, there exists a competitive equilibrium.

Definition 2. *An arrangement (Ψ, p) is a* competitive equilibrium *if for all agents $i \in I$,*

$$\Psi_i \in D_i(p).$$

In other words, a competitive equilibrium is a price vector p and a set of trades Ψ such that each agent demands the set of trades from Ψ he is involved in at the price vector p. The following theorem holds:

Theorem 1. *Let $(I, \Omega, \{u_i\}_{i \in I})$ be a trading network economy such that for all agents $i \in I$, u_i satisfies the FSC condition. Then there exists a competitive equilibrium.*

The proof proceeds in two major steps, which parallel the proof given in Hatfield et al. (2013). First, we construct the reduced economy and show that every agent's utility in this economy satisfies the gross substitutes and complements condition defined in Sun and Yang (2006). Using the equilibrium of the reduced economy, we then construct an arrangement for the original economy and show that it is a competitive equilibrium.

Step 1: Constructing the Reduced Economy

We first briefly describe the kind of economies that are studied in Sun and Yang (2006): There is a set of agents J and a set of items S that can be partitioned into two sets $S = S^1 \sqcup S^2$. Agent $i \in J$ has a valuation function $v_i : 2^S \to \mathbb{R}$ over sets of objects which, together with a price vector $p \in \mathbb{R}^S$, also induces a quasi-linear utility function V_i with

$$V_i(A, p) := v_i(A) - \sum_{a \in A} p_a, \quad A \subseteq S.$$

The demand correspondence in this economy is defined as above (with respect to V_i) and denoted $E_i(p)$.

Definition 3. *The valuation function v_i of agent i satisfies the* gross substitutes and complements (GSC) *condition, if the following holds:*

1. *Fix price vectors $p \leq q$ such that $p_a = q_a$ for all $a \in S^2$. Then for every $A \in E_i(p)$ there exists $B \in E_i(q)$ such that*

$$\begin{aligned} \{a \in A^1 \mid p_a = q_a\} &\subseteq B^1 \quad \text{and} \\ A^2 &\supseteq B^2. \end{aligned} \tag{2}$$

2. *The above condition holds if the sets S^1 and S^2 are exchanged.*

Definition 4. *A* competitive equilibrium *is a price vector p together with a partition $S = X_1 \sqcup \cdots \sqcup X_J$ of S (a set X_i may be empty) and such that*

$$X_i \in E_i(p) \quad \text{for all } i \in J.$$

Now given a trading network economy $(I, \Omega, \{u_i\}_{i \in I})$, we describe how it is transformed into the two-sided economy $(J, S, \{v_i\}_{i \in J})$. First set $J = I$ and $S^1 = \Omega^1, S^2 = \Omega^2$, that is, the set of agents in the two-sided economy is the set of agents in the network economy and the set of objects is the set of trades from the network economy (keeping the distinction between software and hardware trades).

Next, we define the new valuation functions $v_i, i \in J = I$, through

$$v_i(A) := u_i(A_{\to i} \cup (A^c)_{i \to}) - u_i(\Omega_{i \to}) \quad \text{if } A \subseteq \Omega_i,$$

and $v_i(A) = \Pi$ otherwise, where Π is some number low enough such that the agent will never demand such a bundle. Such a number exists since there are only finitely many value the functions u_i can attain, and since by the definition of v_i we have $v_i(\emptyset) = 0$, which is also required in the model of Sun and Yang (2006).

Let us for some set $A \subseteq \Omega_i$ define $A^\# := A_{\to i} \cup (A^c)_{i \to}$. Then it follows that $(A^\#)^\# = A$ and $v_i(A) = u_i(A^\#) - u_i(\Omega_{i \to})$.

Lemma 1. *For each agent $i \in J$, if the valuation functions u_i satisfy property FSC, then the valuation function v_i satisfies property GSC with respect to $S^1 \sqcup S^2$.*

Proof. First note that by the definition of v_i it follows that for any price vector p and any set of trades $A \subseteq \Omega_i$,

$$A \in D_i(p) \quad \Leftrightarrow \quad A^\# \in E_i(p). \tag{3}$$

To see this, we write

$$\begin{aligned} V_i(A^\#, p) &= v_i(A^\#) - p(A^\#) \\ &= u_i(A) - u_i(\Omega_{i \to}) - \big(p(A_{\to i}) + p(\Omega_i) - p(A_{i \to})\big) \\ &= U_i(A, p) + \text{const.} \end{aligned}$$

Hence, A maximizes U_i if and only if $A^\#$ maximizes V_i.

Now fix any two price vectors $p \leq q$ such that $p_a = q_a$ for all $a \in S^2$, and pick a set $A \in E_i(p)$. We have to find a set $B \in E_i(q)$ fulfilling condition (2). Since by (3), $A^\# \in D_i(p)$ we can plug it into the definition of FSC and get a set $B^\# \in D_i(q)$ fulfilling condition (1).

We claim that $B = (B^{\#})^{\#} \in E_j(q)$ is the desired set. Consider $\omega \in \Omega^1_{\to i}$ with $p_\omega = q_\omega$. Then condition (1) together with the definition of the indicator function says

$$\omega \in (A^{\#})^1_{\to i} = A^1_{\to i} \quad \Rightarrow \quad \omega \in (B^{\#})^1_{\to i} = B^1_{\to i}.$$

Similarly, if $\omega \in \Omega^1_{i\to}$, we have

$$\omega \notin (A^{\#})^1_{i\to} = (A^c)^1_{i\to} \quad \Rightarrow \quad \omega \notin (B^{\#})^1_{i\to} = (B^c)^1_{i\to},$$

which translates to $\omega \in A^1_{i\to} \;\Rightarrow\; \omega \in B^1_{i\to}$. Putting both implications together, we have shown that $\{\omega \in A^1 \mid p_\omega = q_\omega\} \subseteq B^1$.

The cases where $\omega \in \Omega^2_i$ are treated analogously, showing that $A^2 \supseteq B^2$, which means that B fulfills condition (2). □

Step 2: Constructing the Equilibrium

Since valuation functions in the reduced economy $(J, S, \{v_i\}_{i\in I})$ satisfy the GSC condition, by Theorem 3.1 in Sun and Yang (2006), there exists a competitive equilibrium with price vector p and partition $S = X_1 \sqcup \cdots \sqcup X_J$. Note that in such an equilibrium, all objects/trades are assigned to some agent.

For some trade ω let $\mu(\omega)$ be the agent who receives the trade in equilibrium, i.e., $\omega \in X_{\mu(\omega)}$. Also, let $b(\omega)$ be the buyer agent of trade ω in the network economy, and let $s(\omega)$ be the seller agent, i.e., $\omega = (s(\omega), b(\omega))$. Note that every trade is assigned to one agent, and by the construction of v_i every trade is either assigned to its buying or its selling agent.

Now construct the arrangement (Ψ, p^*) of the network economy as follows: Set $p^* := p$ and $\Psi := \{\omega \in S = \Omega \mid \mu(\omega) = b(\omega)\}$.

Lemma 2. *The arrangement (Ψ, p^*) constitutes a competitive equilibrium of the trading network economy.*

Proof. By construction of the two-sided market and Ψ, in equilibrium, agent i receives the trades $\Psi^{\#}_i$ in the equilibrium of the two-sided economy, i.e., $X_i = \Psi^{\#}_i$. Since p^* and $\{X_i\}_{i\in I}$ is an equilibrium of the reduced economy, we have $\Psi^{\#}_i \in E_i(p^*)$, implying $\Psi_i \in D_i(p^*)$. Hence, (Ψ, p^*) is an equilibrium of the original economy. □

This completes the proof of Theorem 1.

Relation to Previous Reductions

In this section we aim to explain how the above reduction relates to previous reductions to the Kelso and Crawford (1982) model. This also helps to understand in what way the present chapter unifies both Hatfield et al. (2013) and Sun and Yang (2006).

We begin by explaining the reduction of Hatfield et al. (2013). In their model, all the goods underlying the trades are substitutes for the agents. The only difference to Kelso and Crawford (1982) is the network structure and that agents can be involved in trades in which they are sellers. Therefore, the main purpose of the reduction is to translate the network structure into a two-sided economy where every agent only buys

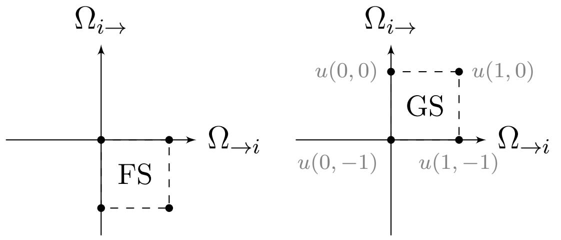

(a) Reduction from full substitutes (FS) to gross substitutes (GS) in Hatfield et al. (2013).

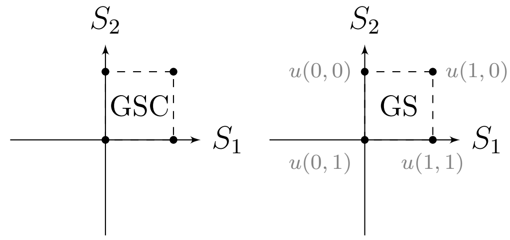

(b) Reduction from gross substitutes and complements to gross substitutes in Sun and Yang (2006).

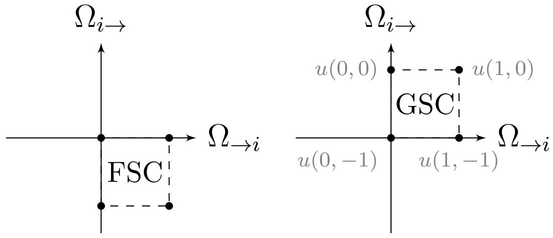

(c) The reduction in this chapter.

Figure 1: Schematic illustration of the reductions. Original valuations (u) are shown on the left hand side, transformed valuations (v) on the right hand side. The gray labels indicate where the function values of v come from.

items. We demonstrate this by writing down the model in an equivalent way, as is suggested by the buy/sell structure and the definition of the indicator function $e(A)$.

We translate every valuation function u_i to a valuation function $u'_i : X_i \to \mathbb{R}$ as follows (cf. Baldwin and Klemperer 2013): The new domain is some subset

$$X_i \subseteq \{x \in \{-1,0,1\}^{\Omega_i} \mid x_\omega = 1 \Rightarrow \omega \in A_{\to i}, x_\omega = -1 \Rightarrow \omega \in A_{i\to}\},$$

and we set $u'_i(x) = u_i(A_i)$ if x can be expressed as $x = e(A_i)$ for some $A_i \subseteq \Omega_i$. If we define $U'_i(x,p) = u'_i(x) - p \cdot x$, then for the respective demand correspondence D'_i it is easy to see that

$$A_i \in D_i(p) \quad \Leftrightarrow \quad e(A_i) \in D'_i(p),$$

so in this sense u_i and u'_i are equivalent. Clearly, for the valuation functions u'_i, an adapted definition of the full substitutes condition can be formulated and the valuation functions for the reduced economy would then be defined as follows (ignoring the normalization term; see Figure 1a):

$$v'_i(x) = u'_i(x - e(\Omega_{i\to}))$$

It is then also clear that v'_i satisfies the gross substitutes condition whenever u'_i satisfies full substitutes, since it is invariant with respect to a shift of the domain.

In Sun and Yang (2006), the economy is reduced to the one in Kelso and Crawford (1982) by essentially "mirroring" the valuation function along the axes of one part of the goods. This is also illustrated in Figure 1b. In contrast to Hatfield et al. (2013), the purpose of this reduction is not to transform a buy/sell economy to a buy-only economy. Also, goods in different sets are genuine complements to each other, whereas in Hatfield et al. (2013) all goods are substitutes. The interpretation that buy and sell trades complement each other is suggested by the network structure.

Analogous to the reduction in Hatfield et al. (2013), the purpose of the reduction in this chapter is to translate the network structure to a two-sided structure: If the model was written in terms of the valuation functions u_i' above, the reduction could be illustrated by Figure 1c. The result is that the goods underlying the trades in the economy can be divided into two parts, where goods in different parts complement each other. Additionally, in the same way the Sun and Yang (2006) model cannot be extended to more than two sets of goods, the network model cannot be extended to more than two sets of items underlying the trades.

Welfare Theorems

Since the welfare theorems are independent of the notion of FSC, the same results as in Hatfield et al. (2013) still hold in our setting. We call a set of trades $\Phi \subseteq \Omega$ *efficient*, if it satisfies

$$\sum_{i\in I} u_i(\Phi) \geq \sum_{i\in I} u_i(A) \qquad \text{for all } A \subseteq \Omega.$$

Theorem 2 (First Welfare Theorem)**.** *Let (Ψ, p) be a competitive equilibrium. Then the set of trades Ψ is efficient.*

Theorem 3 (Second Welfare Theorem)**.** *Suppose the agents' preferences satisfy FSC (so that a competitive equilibrium exists). Then for any competitive equilibrium (Ψ, p) and every efficient set of trades Φ, (Φ, p) is a competitive equilibrium.*

4. Anonymous Prices and Stability

In this section we confirm that all the theorems in Hatfield et al. (2013) concerning the existence of anonymous prices and the relation to stable outcomes carry over to our model. Since the proofs are mostly similar to those in Hatfield et al. (2013), we defer them to the appendix.

Anonymous Prices

In order for our model to generalize the preceding models (Gul and Stacchetti 1999, Sun and Yang 2006, in particular), one needs to establish that under certain circumstances competitive equilibria with anonymous prices exist. For instance, in the model of Gul and Stacchetti (1999), an object (which we would model as an agent) has no "preferences" over which agent it is sold to and can only be sold once. Also, in competitive equilibrium, the object should only be assigned one price, which is the same as saying that all trades with the same object should have the same price in competitive equilibrium. Therefore, we define the following (analogous to Hatfield et al. 2013):

Definition 5. *For agent $i \in I$, the trades in the set $A \subseteq \Omega_i$ are*

1. mutually incompatible, *if* $B \notin D_i(p)$ *for all* p *and for all* $B \subseteq \Omega_i$ *with* $|B \cap A| > 1$,

2. perfect substitutes, *if* $u_i(B \cup \{a\}) = u_i(B \cup \{a'\})$ *for all* $B \subseteq \Omega_i \setminus A$ *and for all* $a, a' \in A$.

Mutual incompatibility says that the agent will never demand a set which contains two or more trades from A. Perfect substitutability says that an agent receives the same marginal utility from any of the items in A. Clearly, we would model an object in an auction with preferences that satisfy precisely mutual incompatibility and perfect substitutability for all trades it can "engage in:" the object can only be sold once and it does not matter for the object to which agent it is sold to.

The following theorem says that there exist competitive equilibria where trades that are mutually incompatible and perfect substitutes for an agent all have the same price.

Theorem 4. *Let* (Ψ, p) *be a competitive equilibrium. Then, if the trades in* $A \subseteq \Omega_i$ *are mutually incompatible and perfect substitutes for agent* $i \in I$, (Ψ, q) *is also a competitive equilibrium, where* q *is defined as*

$$q_a = \begin{cases} \overline{p} = \max_{b \in A_{i\rightarrow}} p_b & \text{if } a \in A_{i\rightarrow}, \\ \underline{p} = \min_{b \in A_{\rightarrow i}} p_b & \text{if } a \in A_{\rightarrow i}, \\ p_a & \text{otherwise.} \end{cases}$$

Stability and Competitive Equilibrium

In this section we show that the relationship between the concept of stability and competitive equilibrium is preserved when generalizing to the full substitutes and complements condition: Every set of trades that is formed in a competitive equilibrium is stable and, if one imposes condition FSC, for every stable set of trades, there exist prices that are competitive equilibrium prices for that set of trades.

In order to define stable sets, we introduce some more notation: a pair $(\omega, p_\omega) \in \Omega \times \mathbb{R}$ is called a *contract* and a set of contracts is called *feasible*, if there are no two contracts consisting of the same trade. A feasible set of contracts Y is called an *outcome*, and for such an outcome we can denote the contained set of trades by A and the corresponding prices by the vector $p \in \mathbb{R}^A$. Then we write

$$U_i(Y) := U_i(A, p)$$

for the utility derived from the set of contracts by agent i.

For an arbitrary set of contracts Y, the *choice correspondence* of agent i is defined as

$$C_i(Y) := \underset{\text{feasible } Z \subseteq Y_i}{\arg\max}\ U_i(Z).$$

We are now ready to define what a stable outcome is.

Definition 6. *The outcome* Y *is* stable, *if it is*

(a) individually rational*:* $Y_i \in C_i(Y)$ *for all* i, *and*

(b) unblocked*: there is no feasible non-empty* blocking set *of contracts Z with $Z \cap Y = \emptyset$ and such that for every agent i that is involved in contracts of Z, $Z_i \subseteq A$ for all choices $A \in C_i(Y \cup Z)$.*

Put differently, stability requires that all contracts are accepted by the agents being part of them and that there is no possibility to propose a new set of different contracts (a blocking set) where all contracts are accepted by all agents being part of the new contracts (and who possibly drop some of the old contracts).

The result that any set of contracts induced by a competitive equilibrium is stable, as well as its proof are independent of the agents' preferences, hence we just state the theorem without repeating the proof.

Theorem 5. *Let Y be the outcome induced by the competitive equilibrium (Ψ, p). Then Y is a stable outcome.*

The converse, however, is not true unless we impose the full substitutes and complements condition:

Theorem 6. *Let condition FSC be satisfied and let Y be a stable outcome with associated trades Ψ and prices r. Then r can be extended to prices p for trades not in Ψ such that (Ψ, p) is a competitive equilibrium.*

5. Discussion

We have shown that the full substitutes condition can be generalized to a condition similar to the gross substitutes and complements condition, still guaranteeing existence of a competitive equilibrium. Competitive equilibria are efficient and every efficient set of trades can be supported by competitive equilibrium prices. Further, if two trades are mutually incompatible and perfect substitutes for one agent, they can be assumed to have the same price in competitive equilibrium. Also, we have shown that the relation between competitive equilibria and stable outcomes carries over to the model with more general preferences: Competitive equilibria are stable, and if the preferences satisfy full substitutes and complements, stable outcomes can be supported by competitive equilibrium prices.

Due to the existence of anonymous prices, the model strictly contains and unifies all the models that it builds upon: Hatfield et al. (2013) can be obtained by setting $\Omega^1 = \Omega$ and $\Omega^2 = \emptyset$. Sun and Yang (2006) can be obtained directly by looking at a bipartite network between the objects and the buyers. Trades between buyers and objects in class 1 and 2 are put in Ω^1 and Ω^2, respectively. Note, that the economy of Sun and Yang (2006) is directly contained in our model and no mapping into a different network economy is needed to prove their result. While the preference assumptions of Kelso and Crawford (1982) are generalized by Sun and Yang (2006), and the market structure is generalized by Hatfield et al. (2013), the present chapter unifies these generalizations.

Also, in the samy way the result of Sun and Yang (2006) does not contradict Theorem 2 in Gul and Stacchetti (1999), our result does not contradict Theorem 7 of Hatfield et al. (2013). The reason for this is that GSC preferences in general do not contain *simple* preferences, which are used to obtain the impossibility result.[3]

[3]For instance, in Example 2 in Hatfield et al. (2013), there is no possibility to partition trades into

Appendix

Proof of Theorem 4

Proof. We have to show that for all agents $j \in I$, $U_j(\Psi_j, q) \geq U_j(B, q)$ holds for all $B \subseteq \Omega_j$. Let $\overline{a}$ and $\underline{a}$ be the trades attaining $\overline{p}$ and $\underline{p}$, respectively. Since the trades in A are mutually incompatible, $|\Psi_j \cap A| \leq 1$. If we have an $a \in \Psi_j \cap A$, assume without loss of generality $a \in A_{i\rightarrow}$ (the other case is treated analogously). Then we know that $p_a = \overline{p}$, since by perfect substitutability

$$u_i(\Psi_i) = u_i(\Psi_i \setminus \{a\} \cup \{\overline{a}\})$$

holds and therefore $p_a < \overline{p}$ would contradict $\Psi_i \in D_i(p)$. It follows that the prices in Ψ_j did not change and therefore

$$U_j(\Psi_j, q) = U_j(\Psi_j, p).$$

Now if $j \neq i$, we can derive for any $B \subseteq \Omega_j$ that

$$U_j(\Psi_j, q) = U_j(\Psi_j, p) \geq U_j(B, p) \geq U_j(B, q),$$

where the last inequality follows from the fact that trades where j buys from i got (weakly) more expensive and trades where j sells to i got (weakly) less expensive.

If, on the other hand, $j = i$, by mutual incompatibility, the only sets to consider are B and $B \cup \{a\}$, where $B \subseteq \Omega_i \setminus A$ and $a \in A$. Since prices in B did not change, we have

$$U_i(\Psi_i, q) = U_i(\Psi_i, p) \geq U_i(B, p) = U_i(B, q).$$

Also, by the definition of $\overline{a}$ and by perfect substitutability,

$$U_i(\Psi_i, q) = U_i(\Psi_i, p) \geq U_i(B \cup \{\overline{a}\}, p) = U_i(B \cup \{\overline{a}\}, q) \geq U_i(B \cup \{a\}, q),$$

which completes the proof. □

Proof of Theorem 6

The proof is adapted from Hatfield et al. (2013) and adjusted to our definition of full substitutes and complements. We first prove two lemmas showing ways of modifying the agents' utility functions while preserving condition FSC.

Lemma 3. *Let u_i satisfy FSC and let Ψ be the set of trades and r be the price vector of outcome Y. Define the valuation function $\hat{u}_i$ on $\Omega \setminus \Psi$ as*

$$\hat{u}_i(\Phi) := \max_{A \subseteq \Psi} \big\{ u_i(\Phi \cup A) + r(A) \big\}.$$

Then $\hat{u}_i$ satisfies FSC.

two sets such that all agents satisfy condition FSC: Since the trades ω and ψ are complements for agent i, $\omega \in \Omega^1$ and $\psi \in \Omega^2$. The preferences of $s(\omega)$ and $s(\psi)$ require $\hat{\omega} \in \Omega^1$ and $\hat{\psi} \in \Omega^2$, which means that agent j cannot have FSC preferences.

Proof. Let $\hat{U}_i$ and $\hat{D}_i$ be the utility function and demand correspondence for the new valuation function $\hat{u}_i$. For some set $S \subseteq \Omega_i$, denote $\hat{S} = S \setminus \Psi$ and let $\hat{S}^*$ be the set attaining the maximum in the definition of $\hat{u}_i(\hat{S})$. Let $\hat{p} \leq \hat{q}$ be price vectors for $\Omega \setminus \Psi$ with $\hat{p}_\omega = \hat{q}_\omega$ for $\omega \in \Omega^2$ and let $\hat{A} \in \hat{D}_i(\hat{p})$. We have to find $\hat{B} \in \hat{D}_i(\hat{q})$ such that (1) holds.

We claim that $A = \hat{A}^* \cup \hat{A} \in D_i(p)$, where p is the combination of the prices r and $\hat{p}$ (similarly for q). Assume on the contrary that there exists S with $U_i(S,p) > U_i(A,p)$. It follows from the definition of $\hat{u}$ that

$$\begin{aligned}\hat{u}_i(\hat{S}) + \hat{p}(\hat{S}) &\geq u_i\big(\hat{S} \cup (S \cap \Psi)\big) + r(S \cap \Psi) + \hat{p}(\hat{S}) \\ &= u_i(S) + p(S) \\ &> u_i(A) + p(A) = \hat{u}_i(\hat{A}) + \hat{p}(\hat{A}),\end{aligned}$$

which contradicts $\hat{A} \in \hat{D}_i(\hat{p})$.

Since u_i satisfies FSC, there exists a set $B \in D_i(q)$ such that (1) holds. We claim that $\hat{B} \in \hat{D}_i(\hat{q})$. Assume that this is not the case, i.e., there exists $\hat{S}$ with $\hat{U}_i(\hat{S},\hat{q}) > \hat{U}_i(\hat{B},\hat{q})$. Then it follows from the definition of $\hat{u}_i$ that

$$\begin{aligned}u_i(\hat{S} \cup S^*) + r(S^*) + \hat{q}(\hat{S}) &> u_i(\hat{B} \cup \hat{B}^*) + r(\hat{B}^*) + \hat{q}(\hat{B}) \\ &\geq u_i(B) + q(B),\end{aligned}$$

which contradicts $B \in D_i(q)$. Now because (1) holds for the sets A and B, it clearly holds for the restrictions $\hat{A}$ and $\hat{B}$. □

Lemma 4. *Let $\hat{u}_i$ satisfy FSC and let $\delta > 0$ be given. Define the valuation function $\tilde{u}_i$ on Ω as*

$$\tilde{u}_i(\Psi) := \hat{u}_i(\Psi) - \delta|\Psi_i|.$$

Then $\tilde{u}_i$ satisfies FSC.

Proof. Denote by $\tilde{U}_i$ and $\tilde{D}_i$ the corresponding utility function and demand correspondence. For an arbitrary price p define $\tilde{p}$ through

$$\tilde{p}_\omega := \begin{cases} p_\omega - \delta & \text{if } \omega \in \Omega_{i\to} \\ p_\omega + \delta & \text{if } \omega \in \Omega_{\to i} \\ p_\omega & \text{otherwise.} \end{cases}$$

Then we have $\tilde{U}_i(\Phi,p) = U_i(\Phi,\tilde{p})$ and therefore

$$A \in D_i(\tilde{p}) \quad \Leftrightarrow \quad A \in \tilde{D}_i(p).$$

The lemma then follows from the observation that if the prices p, q satisfy the premises of Definition 1, then the modified prices $\tilde{p}, \tilde{q}$ do so as well. □

The rest of the proof is directly taken from Hatfield et al. (2013):

Proof of Theorem 6. As in Lemma 3, construct $\hat{u}_i$ for all $i \in I$ and let $(\hat{\Psi}, \hat{p})$ be a competitive equilibrium in the economy with trades $\Omega \setminus \Psi$ and valuation functions

$\hat{u}_i$. If we can choose $\hat{\Psi}$ such that $\hat{\Psi} = \emptyset$ we are done, since then, (Ψ, p) (where p is the combination of r and $\hat{p}$) is a competitive equilibrium in the original economy: $\emptyset \in \hat{D}_i(\hat{p})$ implies that no agent wants to add trades not in Ψ, and individual rationality of Y implies that no agent wants to drop a trade in Ψ.

However, if $\emptyset$ is not a competitive equilibrium, we know by Theorem 2 and 3 that $\hat{\Psi}$ is efficient and $\emptyset$ is not. We can then define

$$\varepsilon := \sum_{i \in I} \hat{u}_i(\hat{\Psi}) - \sum_{i \in I} \hat{u}_i(\emptyset) > 0$$

and set $\delta := \frac{\varepsilon}{2|\Omega|}$. Define $\tilde{u}_i$ as in Lemma 4. By the choice of δ, we know that

$$\sum_{i \in I} \tilde{u}_i(\hat{\Psi}) > \sum_{i \in I} \tilde{u}_i(\emptyset),$$

and therefore that $\emptyset$ is no competitive equilibrium for the valuations $\tilde{u}_i$ either. Hence, a competitive equilibrium $(\tilde{\Psi}, \tilde{p})$ for the new valuations will satisfy $\tilde{\Psi} \neq \emptyset$.

It follows that $\tilde{U}_i(\tilde{\Psi}, \tilde{p}) \geq \tilde{U}_i(\Phi, \tilde{p})$ for $\Phi \subset \tilde{\Psi}$ and by the definition of $\tilde{u}_i$ we then have $\hat{U}_i(\tilde{\Psi}, \tilde{p}) > \hat{U}_i(\Phi, \tilde{p})$. We can now see that $Z := \{(\psi, \tilde{p}_\psi) | \psi \in \tilde{\Psi}\}$ is a blocking set: By looking at the definition of $\hat{u}_i$, for every agent i, every choice $A \in C_i(Y \cup Z)$ has to contain Z_i. This contradicts the fact that Y is stable. $\square$

Chapter 5

Discrete Convex Analysis and Tâtonnement for Economies with Indivisibilities

We apply the theory of Discrete Convex Analysis to economies with indivisibilities in order to derive a simple tâtonnement process for settings where agents have substitutes preferences over heterogeneous goods. Specifically, we reinterpret the price adjustment process discovered by Ausubel (2006) in terms of a steepest descent algorithm for the minimization of discrete convex functions and generalize it to settings that allow agents to be producers and/or consumers of multiple units of goods. We apply our model to the substitutes and complements setting introduced by Sun and Yang (2009) and the trading network economy of Hatfield et al. (2013).

1. Introduction

The idea of a tâtonnement process that tentatively adjusts prices according to current supply and demand was formulated by Walras (1874), and has since then been used to derive ascending iterative auctions and methods for finding market-clearing prices. Walras also realized that a form of substitutability of demand was needed for such a process to work. For divisible goods, a formal description of a tâtonnement process and convergence results were first given by Samuelson (1947) and Arrow and Hurwicz (1958). Similarly as in the first papers that established existence of a market-clearing equilibrium (Arrow and Debreu 1954, McKenzie 1959), these results assume that the preferences which induce demand satisfy certain notions of convexity.

In settings with indivisible goods, the first formal tâtonnement process was an algorithm described by Kelso and Crawford (1982) for the allocation of workers to firms. They realized that sufficient for convergence in these settings is that the agents' demand satisfies the gross substitutes condition. In subsequent papers, different algorithms were developed for the same condition (Gul and Stacchetti 1999, Ausubel 2006, Milgrom and Strulovici 2009). As observed recently by Fujishige and Yang (2003), the demand function of an agent satisfies gross substitution precisely if the valuation function that describes the agent's preferences is $M^\natural$-concave (that is, it belongs to a class of well-behaved discrete concave functions). In this chapter, we make use of this connection in order to apply the theory of Discrete Convex Analysis to the design and analysis of discrete tâtonnement processes, and argue that this connection is the driving force behind the convergence results in discrete settings.

To be more precise, we look at economies where every agent has quasi-linear preferences over the consumption and production of bundles of goods. The economy may initially be endowed with arbitrary quantities of the goods in order to model auction environments, and we assume that individuals are endowed with a sufficient amount of money so that they will never face budget constraints and are able to buy any bundle of goods as they wish. The chapter makes use of the theory of Discrete Convex Analysis in order to interpret the auction developed by Ausubel (2006) as a steepest descent algorithm for the conjugate of the aggregate valuation function. Based on deep results from Discrete Convex Analysis, we can give simple and intuitive proofs for the convergence properties, and further are able to generalize Ausubel's process to more general settings. In particular, we argue that an $M^\natural$-convex function is the appropriate extension of a gross substitutes valuation function to multiple units of goods. This extension has several advantages: First, it implies that the tâtonnement process generates linear prices for multiple units of goods (as in Milgrom and Strulovici 2009). This contrasts the work of Ausubel (2006), who treats every unit of every good as a separate item and therefore generates non-linear prices, as well as unnecessarily increases the complexity of the auction. Second, and more important for the applications we present, it allows for models where agents have preferences over negative amounts of goods (so that some agents may be buyers and/or sellers), as well as models that do not satisfy the assumption of free disposal.

In particular, with the latter two features we can gain the following important insights: First, the double-track adjustment process proposed by Sun and Yang (2009), which works for two classes of goods such that items in the same class are substitutes and items in different classes are complements to each other, can be obtained by just "mirroring" all valuation functions along the axes of the goods in one class. Thereby, our algorithm also generalizes the double-track process to multiple units of goods. Second, our algorithm can immediately be applied to the trading network economy of Hatfield et al. (2013) by recognizing that their full substitutes condition is equivalent to the convexity assumption we use in this chapter.

The interpretation of tâtonnement processes as discrete steepest descent algorithms also has the advantage that statements about the algorithmic complexity of the process can be made. Further, efficient scaling techniques can be applied which yield adjustment processes that only need to state a strongly polynomial number of ask prices, albeit at the cost of monotone convergence.

$M^\natural$-convexity turns out to be the right notion of convexity in the discrete setting because $M^\natural$-convex functions exhibit several properties that are important for the establishment of a discrete price adjustment process. First, the class of $M^\natural$-convex functions is closed under aggregation, which implies that aggregate demand shares the same properties as every agent's demand. Since therefore, aggregate demand is convex,[1] every bundle of goods is demanded at some price, which means that a competitive equilib-

[1] The appropriate notion is convex-extensibility, see Definition 3 below.

rium exists. The notion of $M^\natural$-convexity is, however, too strong to be necessary for the existence of competitive equilibria. We refer to Danilov, Koshevoy and Murota (2001) and Baldwin and Klemperer (2013) for a complete characterization of the classes of valuation functions that guarantee existence of an equilibrium. $M^\natural$-convex functions provide, additionally, the appropriate combinatorial properties such as convexity and submodularity of the aggregate indirect utility function which are key for the (monotone) convergence of discrete steepest descent algorithms.

Another key property that is implied by $M^\natural$-convexity is that the optimal descent direction of the indirect utility function at a given price is entirely determined by the aggregate demand correspondence. Therefore, even though indirect utility cannot be elicited directly, the steepest descent algorithm is economically suitable because it can proceed through best response information from the agents. While the algorithm in this chapter guarantees convergence to an equilibrium if the agents reveal their demand truthfully, it does not incentivize them to do so. The algorithm is therefore suitable for settings where agents can be assumed to have no market power and therefore to act as price takers (Klemperer 2010), but can also be seen as an attempt to understand how natural adjustment dynamics work in settings with indivisible goods.

Steepest descent methods for the design of iterative auctions are complemented by primal-dual and linear programming algorithms (Demange, Gale and Sotomayor 1986, Gul and Stacchetti 1999, Parkes and Ungar 2000, deVries, Schummer and Vohra 2007) and related to algorithms for finding stable outcomes in matching models (Gale and Shapley 1962, Ostrovski 2008).

Illustration: Convexity and Tâtonnement

The role of convexity and conjugacy in the establishment of a price adjustment process can best be described visually by means of a valuation function with a continuous domain. In order to keep things simple, we assume that the aggregate valuation function v of a group of agents over quantities x of a single good is given (see Figure 1). Aggregate utility is quasi-linear, and therefore a competitive equilibrium price for endowment $\bar{x}$ is a price p^* such that $\bar{x}$ maximizes the expression $v(x) - p^*x$, i.e., a price such that the agents demand a quantity of precisely $\bar{x}$. Alternatively, the maximization problem of the agents can be written as maximizing the linear function $(-p, 1)^T(x, y)$ over points (x, y) such that $y \leq v(x)$. Since v is concave, it is then clear that for an equilibrium price p^*, the vector $(-p^*, 1)$ will be perpendicular to the tangent of the convex set $V = \{(x, y) \mid y \leq v(x)\}$ at the point $(\bar{x}, v(\bar{x}))$. It is then also clear that concavity of v guarantees that an equilibrium price exists for every endowment.

In order to derive a process that converges to the equilibrium price we look at the indirect utility at price p, which is defined as $U(p) = \max_x v(x) - px$. If x^* attains the maximum, we can write $v(x^*) = U(p) + px^*$ and therefore $U(p)$ is the intercept of the tangent of the set V at the point $(x^*, v(x^*))$. Also, since v is concave, we can see from Figure 1 that the value $v(x)$ can be recovered from the intercepts $U(p)$ via $v(x) = \min_p U(p) + px$. This is called conjugacy and the functions v and U are said to be conjugate to each other. Since U is defined as the maximum over a family of linear functions, it will be convex.

If p^* is an equilibrium price for endowment $\bar{x}$, then $v(\bar{x}) - p^*\bar{x} = U(p^*)$ and therefore,

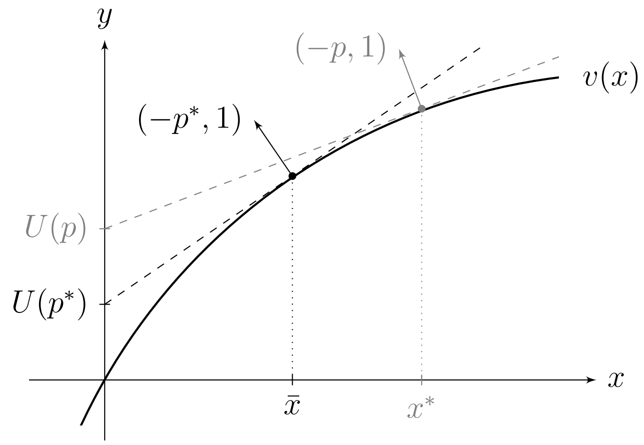

Figure 1: Concavity and conjugacy.

because of conjugacy, p^* attains $\min_p U(p) + p\bar{x}$. Equilibrium prices can therefore be found by computing a minimizer of the function $h(p) = U(p) + p\bar{x}$. Since U is convex, h is convex as well, and there are well-developed algorithms for minimizing convex functions. Some examples are steepest descent and gradient methods.

This chapter demonstrates that these considerations can be applied to valuation functions with discrete domains as well by using the theory of Discrete Convex Analysis. We proceed as follows: Section 2 introduces the notation and basic model. Different restrictions on preferences that are equivalent to gross substitutes are discussed in Section 3. Section 4 provides a brief introduction to Discrete Convex Analysis. Then, Section 5 covers the existence and properties of competitive equilibria. Section 6 presents the price adjustment process. Finally, in Section 7, we apply our results and conclude in Section 8.

2. Basic Model

In this section we introduce the notation that is used throughout the chapter and present the basic economy. We also define the concept of competitive equilibrium.

Notation

Let E be some finite ground set and consider the vector space $\mathbb{R}^E$. For some vector $x \in \mathbb{R}^E$ and element $e \in E$, we write $x_e = x(e)$ for the eth component of x. By $x_{-e} \in \mathbb{R}^{E \setminus \{e\}}$ we mean the vector x without component e. We define the *characteristic vector* $\mathbf{1}_S \in \mathbb{R}^E$ of $S \subseteq E$ as

$$\mathbf{1}_S(e) = \begin{cases} 1 & \text{if } e \in S \\ 0 & \text{otherwise.} \end{cases}$$

For an element $e \in E$, we simply write $\mathbf{1}_e = \mathbf{1}_{\{e\}}$ and by convention we write $\mathbf{1}_0 = (0, \ldots, 0)$. The maximum norm of a vector $x \in \mathbb{R}^E$ is defined as $||x||_\infty = \max_{e \in E} |x(e)|$. For two vectors $x, y \in \mathbb{R}^E$, $x \leq y$ means $x(e) \leq y(e)$ for all $e \in E$, and $x < y$ means

$x \leq y$ and $x(e) < y(e)$ for some $e \in E$. Given some subset $S \subseteq \mathbb{R}^E$, the *indicator function* δ_S is defined as

$$\delta_S(x) = \begin{cases} 0 & \text{if } x \in S \\ \infty & \text{otherwise.} \end{cases}$$

We define the *positive and negative support* of a vector $x \in \mathbb{R}^E$ as

$$\text{supp}^+(x) = \{e \in E \mid x(e) > 0\} \qquad \text{and} \qquad \text{supp}^-(x) = \{e \in E \mid x(e) < 0\}.$$

The *inner product* $\langle \cdot, \cdot \rangle : \mathbb{R}^E \times \mathbb{R}^E \to \mathbb{R}$ is defined through

$$\langle x, y \rangle = \sum_{e \in E} x_e y_e.$$

Given a function $f : \mathbb{Z}^E \to \mathbb{R} \cup \{-\infty\}$, the *effective domain* is $\text{dom} f = \{z \in \mathbb{Z}^E \mid f(z) \neq -\infty\}$. The effective domain is defined similarly for functions $f : \mathbb{Z}^E \to \mathbb{R} \cup \{\infty\}$.

Economy and Competitive Equilibrium

There is a set of agents N and a set of *heterogeneous goods* G. The economy is endowed with positive or negative[2] quantities of goods $\bar{x} \in \mathbb{Z}^G$. Typically, in an auction setting, $\bar{x}$ will be the vector of goods for sale and in a pure exchange economy, we will typically have $\bar{x} = 0$. Each agent $i \in N$ has a *valuation function* $v_i : \mathbb{Z}^G \to \mathbb{Z} \cup \{-\infty\}$ over bundles of goods (where multiple units of a good are allowed). The interpretation of $v_i(x) = -\infty$ is that bundle x is infeasible. We require that $\text{dom}\, v_i$ is finite; this is a technical assumption which is implicitly satisfied in all settings without producers (auction environments) and can also be assumed if it is infeasible for producers to produce infinite quantities of a good. Given linear prices $p \in \mathbb{Z}^G$, an agent i derives *quasi-linear utility* $u_i(x, p) = v_i(x) - \langle p, x \rangle$ from bundle $x \in \mathbb{Z}^G$. Agent i's *indirect utility function* is defined as

$$U_i(p) = \max_{x \in \mathbb{Z}^G} v_i(x) - \langle p, x \rangle.$$

We assume that the agents are endowed with a sufficiently high amount of the divisible good, money, so that they never run into budget problems and can freely choose their most-preferred bundle at the stated prices. The *demand correspondence* $D_i(p)$ is the set of maximizers in the expression above.

Definition 1. *A* competitive equilibrium *is an allocation of goods $x^i \in \mathbb{Z}^G$, $i \in N$ with $\sum_{i \in N} x^i = \bar{x}$, together with a price vector $p^* \in \mathbb{R}^G$ such that $x^i \in D_i(p^*)$ for every agent $i \in N$.*

[2]Negative endowments are important for the application to the double-track adjustment process by Sun and Yang (2009), see Section 7.

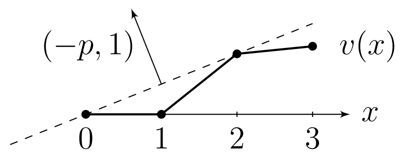

Figure 2: Non-existence of competitive equilibria: There is no price such that the agent demands a quantity of one.

3. Preferences and Discrete Concavity

In settings where only one unit of every good is available and where there is no production, the gross substitutes property has turned out to be sufficient for the existence of competitive equilibria and the design of price adjustment processes. However, when multiple units of a good are available, this condition is too weak. In this section we motivate why $M^\natural$-concavity is a sensible generalization of the gross substitutes condition to our setting. The following is a definition of the gross substitutes property, naively adapted to multiple units of goods (sometimes also called "weak" or "ordinary" substitutes):

Definition 2. *The valuation v_i satisfies* weak/ordinary substitutes (wGS)[3] *if for every pair of prices $p \leq p'$, and every $x \in D_i(p)$, there exists some $x' \in D_i(p')$ such that $x_j \leq x'_j$ for every j for which $p_j = p'_j$.*

To see why this condition is too weak for the existence of a competitive equilibrium, note that it is vacuous for the case of only one good. See the example in Figure 2 for an illustration of non-existence of an equilibrium (one good, one agent). The figure suggests that a form of concavity is required for the existence of a competitive equilibrium. Our requirement on valuation functions (Assumption 1 below) implies that the function is concave in the following sense. Also see Figures 3a and 3b for an illustration.

Definition 3. *A valuation function v_i is* concave-extensible *if it coincides with its concave closure*

$$\bar{v}_i(x) = \inf_{p \in \mathbb{R}^G, \alpha \in \mathbb{R}} \{\langle p, x\rangle + \alpha \mid \langle p, x\rangle + \alpha \geq v_i(z)\ \forall z \in \mathbb{Z}^G\}$$

on the set of integer vectors, i.e., if $v_i(x) = \bar{v}_i(x)$ for all $x \in \mathbb{Z}^G$.

When restricted to the unit cube $\{0,1\}^G$, there are several properties that are equivalent to gross substitutes: For instance, a valuation function satisfies the gross substitutes condition if and only if it satisfies the step-wise gross substitutes condition (Danilov, Koshevoy and Lang 2003):

Definition 4. *Valuation function v_i satisfies* step-wise gross substitutes (SWGS) *if for any $p \in \mathbb{R}^G, x \in D_i(p)$ and $j \in G$, we either have*

[3]Historically, gross substitution is a condition on the demand correspondence of an agent. However, through the specific definition of D_i above, it can be defined in terms of the valuation function v_i. The same applies to the other definitions given in this section.

(i) $x \in D_i(p + \delta \mathbf{1}_j)$ *for all* $\delta \geq 0$ *or*

(ii) there is some $\delta \geq 0$ *and* $x' \in D_i(p + \delta \mathbf{1}_j)$ *with* $x'_j = x_j - 1$ *and* $x'_{-j} \geq x_{-j}$.

The second property which is on the unit cube equivalent to gross substitutes is the single-improvement property (Gul and Stacchetti 1999):

Definition 5. *The valuation function* v_i *satisfies the* single-improvement property*, if for every price* p*, and bundle* $x \notin D_i(p)$*, there exists a bundle* x' *such that* $u_i(x', p) > u_i(x, p)$ *and* $x' = x + \mathbf{1}_j - \mathbf{1}_k$ *with* $j \in \operatorname{supp}^-(x - x') \cup 0$ *and* $k \in \operatorname{supp}^+(x - x') \cup 0$.

Murota and Tamura (2002) show that for general valuation domains, a valuation function satisfies the single-improvement property if and only if it is concave-extensible and satisfies step-wise gross substitutes. Therefore, these properties are suitable generalizations of gross substitutes to multiple units of goods.

Assumption 1. *For every agent* i*, the valuation function* v_i *is concave-extensible and satisfies the step-wise gross-substitutes property. Equivalently,* v_i *satisfies the single-improvement property.*

The interpretation that an agent with a valuation function satisfying this requirement views the different goods as substitutes remains valid also for multiple units. First, on the unit cube, valuations satisfying Assumption 1 are precisely those that satisfy the gross substitutes property as defined by Kelso and Crawford (1982). Second, for non-negative goods vectors, these valuations are precisely the strong substitutes valuations as defined in Milgrom and Strulovici (2009). A valuation satisfies *strong substitutes* if, when every unit of every good is treated as a separate good, the valuation satisfies gross substitutes with respect to them.

It is also known (Fujishige and Yang 2003, Murota and Tamura 2002) that a function satisfies Assumption 1 if and only if it is $M^\natural$-concave. These functions form a class of well-behaved concave functions that play an important role in Discrete Convex Analysis and exhibit combinatorial properties that allow the design of an iterative tâtonnement process. $M^\natural$-concave functions are introduced in the next section. The reader already familiar with Discrete Convex Analysis may want to skip that section.

4. Discrete Convex Analysis

In this section we review the main definitions and results from Discrete Convex Analysis. For a complete and self-contained treatment of the topic we refer the reader to Murota (2003).

A convex function f with a convex domain in $\mathbb{R}^E$ has several attractive properties. First, local optimality implies global optimality. This yields many efficient optimization methods for convex functions. Second, by the supporting hyperplane theorem, the subdifferential of a convex function is non-empty everywhere, and the function can be recovered from the set of subdifferentials. This implies conjugacy and duality results for the convex conjugate (or Legendre-Fenchel transform) of f. The theory of Discrete Convex Analysis identifies classes of convex functions defined on a subset of the discrete

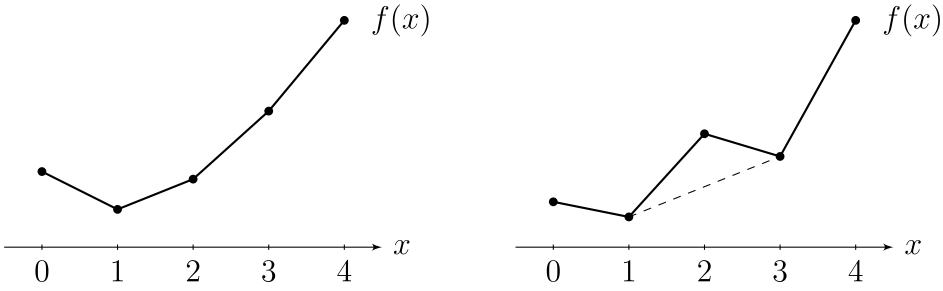

(a) An integrally convex function with local convex extension.

(b) A function that is not integrally convex, together with its local convex extension and convex closure.

Figure 3: Local convex extension, convex closure and integral convexity.

lattice $\mathbb{Z}^E$ for which discrete analogues of the above properties hold. These will play an important role in establishing the results in the remainder of this chapter.

The first important property of functions $f : \mathbb{Z}^E \to \mathbb{R}$, which will be shared by the two subclasses of $M^\natural$- and $L^\natural$-convex functions, is integral convexity. It is defined in terms of suitable convex extensions of f to a real-valued domain. Define the *convex closure* $\bar{f}$ of f as

$$\bar{f}(x) = \sup_{p \in \mathbb{R}^E, \alpha \in \mathbb{R}} \{\langle p, x\rangle + \alpha \mid \langle p, x\rangle + \alpha \leq f(y) \ \forall y \in \mathbb{Z}^E\}.$$

This is equivalent to taking the convex hull of the epigraph of f. If the convex closure coincides with f on the set of integer vectors, i.e., if $f(x) = \bar{f}(x)$ for all $x \in \mathbb{Z}^E$, f is called *convex-extensible*, see Definition 3. We can relax the requirement in the above definition to obtain a local version of the convex extension: The *integral neighborhood* $N(x)$ of $x \in \mathbb{R}^E$ is defined as

$$N(x) = \{y \in \mathbb{Z}^E \ : \ ||y - x||_\infty < 1\}.$$

If we only impose the inequality $\langle p, x\rangle + \alpha \leq f(y)$ for points y in the integral neighborhood of x, we get the *local convex extension* $\tilde{f}$ of f, which is defined as

$$\begin{aligned}\tilde{f}(x) &= \sup_{p \in \mathbb{R}^E, \alpha \in \mathbb{R}} \{\langle p, x\rangle + \alpha \mid \langle p, x\rangle + \alpha \leq f(y) \ \forall y \in N(x)\} \\ &= \inf \left\{ \sum_{y \in N(x)} \lambda_y f(y) \ \middle| \ \sum_{y \in N(x)} \lambda_y y = x, \sum_{y \in N(x)} \lambda_y = 1, \lambda_y \geq 0 \right\}. \qquad (1)\end{aligned}$$

Here, equality of the two expressions follows from linear programming duality (see, e.g., Schrijver 1986).

Definition 6. *A function $f : \mathbb{Z}^E \to \mathbb{R}$ is called* integrally convex*, if the local convex* *extension* $\tilde{f}$ *is convex.*[4] *Equivalently,* f *is integrally convex, if* $\tilde{f} = \bar{f}$.

The function f *is called* integrally concave*, if the function* $-f$ *is integrally convex.*

A set $S \subseteq \mathbb{Z}^E$ *is an* integrally convex set*, if its indicator function* δ_S *is integrally convex.*

Integrally convex functions share with convex functions the important property that local minima are also global minima.

Proposition 1. *Let* $f : \mathbb{Z}^E \to \mathbb{R}$ *be an integrally convex function and* $x \in \mathbb{Z}^E$*. Then* $f(x) \leq f(y)$ *for all* $y \in \mathbb{Z}^E$ *if and only if* $f(x) \leq f(y)$ *for all* $y \in \mathbb{Z}^E$ *with* $||y - x||_\infty \leq 1$.

Proof. We only need to show sufficiency. Consider the local convex extension $\tilde{f}$ of f. From the definition of $\tilde{f}$ in (1) and the local optimality of x with respect to f it follows that $\tilde{f}(x) \leq \tilde{f}(y)$ for all y with $||y - x||_\infty \leq 1$. Hence, x is a local minimum of $\tilde{f}$, which is convex because of integral convexity of f. Therefore, x is also a global minimum of $\tilde{f}$ and in particular of f. □

While integral convexity is sufficient for the global optimality of local minima, more combinatorial structure is needed for the conjugacy and duality results we need. Discrete Convex Analysis identifies $M^\natural$-convex and $L^\natural$-convex functions as two important classes of integrally convex functions which are in one-to-one correspondence to each other under the Legendre-Fenchel transformation.

Definition 7. *A function* f *is* $M^\natural$-convex*, if for* $x, y \in \operatorname{dom} f$ *and* $j \in \operatorname{supp}^+(x - y)$

(i) $f(x) + f(y) \geq f(x - \mathbf{1}_j) + f(y + \mathbf{1}_j)$ *or*

(ii) $f(x) + f(y) \geq f(x - \mathbf{1}_j + \mathbf{1}_k) + f(y + \mathbf{1}_j - \mathbf{1}_k)$ *for some* $k \in \operatorname{supp}^-(x - y)$.

A function f *is* $M^\natural$-concave *if the function* $-f$ *is* $M^\natural$*-convex.*

A set $X \subseteq \mathbb{Z}^E$ *is an* $M^\natural$-convex set*, if its indicator function* δ_X *is* $M^\natural$*-convex.*

The exchange property *(ii)* is closely related to the exchange axiom in matroid theory. Therefore, the M stands for "matroid."

Definition 8. *A function* g *is* $L^\natural$-convex*, if for all* $p, q \in \mathbb{Z}^G$ *and all* $\alpha \in \mathbb{Z}_+$,

$$g(p) + g(q) \geq g([p - \alpha\mathbf{1}] \vee q) + g(p \wedge [q + \alpha\mathbf{1}]).$$

A function g *is* $L^\natural$-concave *if the function* $-g$ *is* $L^\natural$*-convex.*

A set $P \subseteq \mathbb{Z}^E$ *is an* $L^\natural$-convex set*, if its indicator function* δ_P *is* $L^\natural$*-convex.*

$L^\natural$-convex functions are precisely those submodular functions that are integrally convex. The submodularity of an $L^\natural$-convex function can be obtained by setting $\alpha = 0$ in the definition above. Since $L^\natural$-convex sets are precisely the integral sublattices of Z^E, the L stands for "lattice."

[4]Note that integrally convex functions are convex-extensible, but the reverse is not necessarily true. See Example 3.20 in Murota (2003).

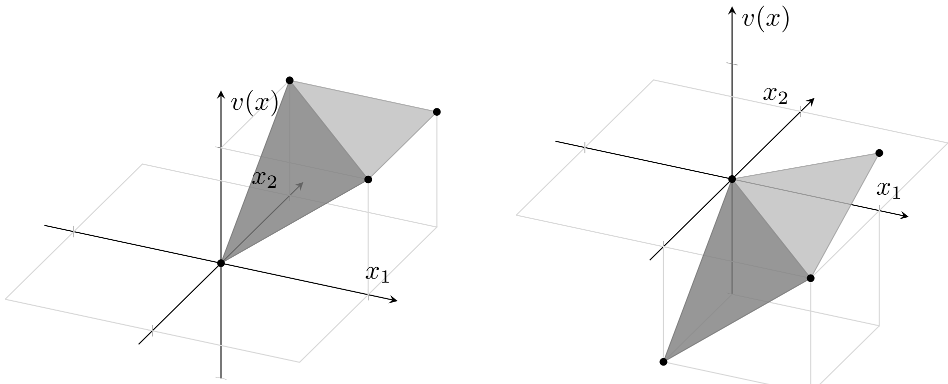

(a) An $M^\natural$-concave function f. **(b)** An $L^\natural$-convex function g.

Figure 4: An integrally concave and convex function. Note that the sets $\arg\max_x f(x) - \langle p, x\rangle$ and $\arg\min_x g(p) - \langle p, x\rangle$ are $M^\natural$-convex and $L^\natural$-convex sets, respectively.

The classes of $M^\natural$- and $L^\natural$-convex functions and sets are in many ways dual to each other. First, the Legendre-Fenchel transform of an $M^\natural$-convex function is $L^\natural$-convex and vice versa. Second, the superdifferential of an $M^\natural$-concave function is an $L^\natural$-convex set, while the subdifferential of an $L^\natural$-convex function is an $M^\natural$-convex set. The superdifferential $\partial f(x)$ of an integrally concave function $f : \mathbb{Z}^E \to \mathbb{R} \cup \{\infty\}$ at $x \in \operatorname{dom} f$ is defined as

$$\partial f(x) = \{p \in \mathbb{Z}^E \mid f(y) - f(x) \leq \langle p, y - x\rangle \ \forall y \in \mathbb{Z}^E\}.$$

Theorem 1. *Let f be an $M^\natural$-concave function with* $\operatorname{dom} f$ *finite. Then*

(i) for all $x \in \operatorname{dom} f$, the superdifferential $\partial f(x)$ is a non-empty $L^\natural$-convex set.

(ii) for all $p \in \mathbb{Z}^E$, the set of maximizers $\arg\max_x\{f(x) - \langle p, x\rangle\}$ *is a non-empty $M^\natural$-convex set.*

An analogue version of the above theorem holds for an $M^\natural$-convex function and its subdifferential. The set of maximizers in part *(ii)* of the theorem are depicted in Figure 4a on the x_1, x_2-plane. It turns out that $M^\natural$-concave functions are characterized by either of the properties in Theorem 1.

Define the subdifferential $\partial g(x)$ of an integrally convex function $g : \mathbb{Z}^E \to \mathbb{R} \cup \{\infty\}$ at $x \in \operatorname{dom} g$ as

$$\partial g(p) = \{x \in \mathbb{Z}^E \mid g(q) - g(p) \geq \langle x, q - p\rangle \ \forall q \in \mathbb{Z}^E\}.$$

Theorem 2. *Let g be an $L^\natural$-convex function with* $\operatorname{dom} g$ *finite. Then*

(i) for all $p \in \operatorname{dom} g$, the subdifferential $\partial g(p)$ is a non-empty $M^\natural$-convex set.

(ii) for all $x \in \mathbb{Z}^E$, the set of minimizers $\arg\min_p\{g(p) - \langle p, x\rangle\}$ *is a non-empty $L^\natural$-convex set.*

An analogue version of this theorem holds for an $L^\natural$-concave function and its superdifferential. The set of $L^\natural$-convex functions is characterized by either of the properties in Theorem 2.

The *concave conjugate or Legendre-Fenchel transform* of a function f is defined as

$$f^\circ(p) = \inf_x \{\langle p, x\rangle - f(x)\}.$$

The following conjugacy result holds for an $M^\natural$-concave function (an analogue result holds for $L^\natural$-convex functions).

Theorem 3 (Conjugacy)**.** *Let f be an $M^\natural$-concave function. Then*

(i) the concave conjugate f° is $L^\natural$-concave and

(ii) the biconjugate of f is identical to f itself: $f^{\circ\circ} = f$.

A special property of $M^\natural$-concavity is that it is preserved under the following operation (Murota 2003). Define the *convolution* $f_1 \,\Box\, f_2$ of two $M^\natural$-concave functions f_1 and f_2 through

$$(f_1 \,\Box\, f_2)(x) = \max_{x^1, x^2} \left\{ f_1(x^1) + f_2(x^2) \mid x^1 + x^2 = x \right\}.$$

We can see from the definition that the effective domain of $f_1 \,\Box\, f_2$ will be the *Minkowski sum* $\operatorname{dom} f_1 + \operatorname{dom} f_2 = \{x = x^1 + x^2 \in \mathbb{R}^E \mid x^1 \in \operatorname{dom} f_1 \text{ and } x^2 \in \operatorname{dom} f_2\}$. Since the convolution is associative, we can for a collection of functions $\{f_i\}_{i\in N}$, define $f_N = \Box_{i\in N} f_i$. It follows that f_N will be $M^\natural$-concave given that f_i is $M^\natural$-concave for all $i \in N$.

The Legendre-Fenchel transformation and the convolution operator satisfy the relation

$$(f_1 \,\Box\, f_2)^\circ = f_1^\circ + f_2^\circ. \tag{2}$$

5. Competitive Equilibrium

This section is concerned with the existence and properties of competitive equilibria. As we show below, existence of competitive equilibria is a property of the aggregate valuation function: If it is concave in the sense that the subdifferential is always non-empty, then an equilibrium exists for every endowment (also see Baldwin and Klemperer 2013, Danilov et al. 2001). This is always fulfilled for integrally convex functions and $M^\natural$-concave functions in particular. For the case of free disposal, we also show that equilibrium prices will always be non-negative.

The definition of the indirect utility function $U_i(p)$ and the concave conjugate v_i° above implies $U_i = -v_i^\circ$. Let v_N denote the aggregate valuation function of all agents in N, i.e.,

$$v_N(x) = \max_{\{x^i\}_{i\in N}} \left\{ \sum_{i\in N} v_i(x^i) \,\middle|\, \sum_{i\in N} x^i = x \right\}.$$

This aggregate valuation function is just the convolution of all the agents' valuation functions. Then, writing U_N for the indirect utility function of the whole group of agents, we get

$$U_N = -v_N^\circ = -\sum_{i\in N} v_i^\circ = \sum_{i\in N} U_i \tag{3}$$

by (2). As before, the aggregate demand of all agents $D_N(p)$ is defined as the set of maximizers attaining $U_N(p)$. This allows us to rephrase the definition of competitive equilibrium.

Proposition 2. *Price vector p^* is a competitive equilibrium price vector if and only if $\bar{x} \in D_N(p^*)$.*

Proof. First, let $x^i \in D_i(p^*)$ and $\sum_{i\in N} x^i = \bar{x}$. Then $U_i(p^*) = v_i(x^i) - \langle p^*, x^i\rangle \geq v_i(x^i) - \langle p^*, y^i\rangle$ for all allocations $\{y^i\}_{i\in N}$, $\sum_{i\in N} y^i = \bar{x}$. Summing up these inequalities implies $\sum_{i\in N} v_i(x^i) \geq \sum_{i\in N} v_i(y^i)$ (cf. First Welfare Theorem) and therefore $v_N(\bar{x}) = \sum_{i\in N} v_i(x^i)$.

Using (3), we have

$$U_N(p^*) = \sum_{i\in N}[v_i(x^i) - \langle p^*, x^i\rangle] = v_N(\bar{x}) - \langle p^*, \bar{x}\rangle,$$

which implies $\bar{x} \in D_N(p^*)$.

Conversely, let $\bar{x} \in D_N(p^*)$. We get from the definition of v_N an allocation $\bar{x} = \sum_{i\in N} x^i$ with $v_N(\bar{x}) = \sum_{i\in N} v_i(x^i)$. Then we get

$$\sum_{i\in N} U_i(p^*) = U_N(p^*) = v_N(\bar{x}) - \langle p^*, \bar{x}\rangle = \sum_{i\in N}[v_i(x^i) - \langle p^*, x^i\rangle],$$

where we in turn use (3), $\bar{x} \in D_N(p^*)$, and the definition of v_N. Since $U_i(p^*) \geq v_i(x^i) - \langle p^*, x^i\rangle$ for all $i \in N$, these inequalities need to hold with equality and therefore $x^i \in D_i(p^*)$ for all $i \in N$. □

The proof also implies that $D_N(p)$ is equal to the Minkowski sum $\sum_{i\in N} D_i(p)$, which is defined as $A + B = \{a + b \mid a \in A \text{ and } b \in B\}$. The proposition indicates that the existence of competitive equilibria is a property of the aggregate valuation function v_N. If it is concave in the sense that for every endowment $\bar{x}$, the subdifferential at $\bar{x}$ is non-empty, then a competitive equilibrium is guaranteed to exist. Since v_N is $M^\natural$-concave, this is the case in our setting.

Theorem 4. *In the economy defined above, if the valuation function of every agent $i \in N$ satisfies Assumption 1 and $\bar{x} \in \operatorname{dom} v_N$, a competitive equilibrium exists.*

Proof. Since v_i is $M^\natural$-concave for all $i \in N$, the aggregate valuation function v_N is also $M^\natural$-concave. Also note that for $p \in \mathbb{Z}^G$ and $x \in \operatorname{dom} v_N$, we have $x \in D_N(p) \Leftrightarrow p \in \partial v_N(x)$. By Theorem 1 (i), for $M^\natural$-concave functions $\partial v_N(\bar{x})$ is non-empty and therefore there exists p^* such that $\bar{x} \in D_N(p^*)$, that is, a competitive equilibrium exists. □

Remark 1. *Theorem 1 also implies that the set of competitive equilibrium prices is an $L^\natural$-convex set. In particular this means that it is a lattice.*

Although we do not need to impose free disposal for our results, the following proposition establishes the intuitive fact that, whenever there is free disposal, prices are non-negative in equilibrium.

Proposition 3. *Assume that v_i is non-decreasing for all $i \in N$, that is $v_i(x) \leq v_i(y)$ for $x \leq y$. Then for every competitive equilibrium price p^* we have $p^* \geq 0$.*

Proof. We first show that if v_1 and v_2 are non-decreasing, then $v_1 \,\square\, v_2$ is non-decreasing. Then by induction it follows that v_N is non-deceasing. So let $x \leq y$ and $x^1 + x^2 = x$ such that $(v_1 \,\square\, v_2)(x) = v_1(x^1) + v_2(x^2)$. Since v_2 is non-decreasing, we have

$$v_1(x^1) + v_2(x^2) \leq v_1(x^1) + v_2(x^2 + [y - x]).$$

Then $(v_1 \,\square\, v_2)(x) \leq (v_1 \,\square\, v_2)(y)$ follows because $y = x^1 + x^2 + [y - x]$.

Now take $j \in G$ and define $y = \bar{x} + \mathbf{1}_j$. Since $p^* \in \partial v_N(\bar{x})$, we have

$$v_N(y) - v_N(\bar{x}) \leq \langle p^*, y - \bar{x} \rangle = p_j^*$$

and therefore $p_j^* \geq 0$, which completes the proof. □

6. Tâtonnement

Conjugacy between the aggregate valuation and indirect utility function helps us to easily see that the set of competitive equilibrium prices coincides with the set of minimizers of a certain function (Ausubel 2006, Milgrom and Strulovici 2009). Since this function will be $L^\natural$-convex, steepest descent algorithms can be used for the computation of competitive equilibrium prices. In this section, we present such an algorithm and analyze its convergence properties. We also make use of the fact that the subdifferential of an $L^\natural$-convex function determines its slope in certain directions to show how one can compute the descent direction via the demand sets.

Proposition 4. *Price vector p^* supports endowment $\bar{x}$ in competitive equilibrium if and only if it minimizes the function $h(p) = \langle p, \bar{x} \rangle + U_N(p)$.*

Proof. By Proposition 2, a price vector p^* supports $\bar{x}$ in competitive equilibrium if and only if

$$v_N(\bar{x}) - \langle p^*, \bar{x} \rangle = U_N(p^*) = -v_N^\circ(p^*).$$

By Theorem 3 (conjugacy) we know that $v_N = v_N^{\circ\circ}$ and therefore $v_N(\bar{x}) = \inf_p \{\langle p, \bar{x} \rangle - v_N^\circ(p)\}$. Hence, necessity follows since p^* attains the infimum and therefore minimizes h. Conversely, if p^* minimizes h, it attains the infimum and therefore constitutes an equilibrium price vector. □

For general valuation functions, a price that minimizes h is often referred to as a *quasi-equilibrium* (Milgrom and Strulovici 2009). Theorem 3 implies that the function h is $L^\natural$-convex: It is the sum of a linear and an $L^\natural$-convex function. In the following we will present a steepest descent algorithm that minimizes the function h and therefore constitutes a price adjustment process for the economy considered. The correctness of the algorithm follows from the following optimality criterion for the minimization of $L^\natural$-convex functions:

Proposition 5. *For an $L^\natural$-convex function h we have that $h(p) \leq h(q)$ for all $q \in \mathbb{Z}^G$ if and only if $h(p) \leq h(p \pm \mathbf{1}_S)$ for all $S \subseteq G$.*

Proof. We only need to show sufficiency. With regard to Proposition 1 we need to show that $h(q) \geq h(p)$ for all $q \in \mathbb{Z}^G$ with $||q - p||_\infty \leq 1$. Every such q can be written as $q = p + \mathbf{1}_X - \mathbf{1}_Y$ for suitable disjoint $X, Y \subseteq G$. Then submodularity and local optimality of p imply

$$h(p) + h(p + \mathbf{1}_X - \mathbf{1}_Y) \geq h(p + \mathbf{1}_X) + h(p - \mathbf{1}_Y) \geq 2h(p),$$

which completes the proof. □

It is therefore straightforward to use the following algorithm for minimizing h: Start with an arbitrary price vector p and search for a subset of goods S and $\varepsilon \in \{-1, 1\}$ such that $h(p+\varepsilon\mathbf{1}_S) - h(p)$ is minimal. If no subset S and ε can be found such that this difference is negative, p is a competitive equilibrium. Otherwise update the price to $p + \varepsilon\mathbf{1}_S$ and iterate. Since we have integer valuations, h decreases by at least 1 in every step and therefore (since a competitive equilibrium exists) the algorithm converges after finitely many steps.[5] It is summarized in Algorithm 1.

Algorithm 1 Steepest Descent

1. Pick an arbitrary price vector $p \in \operatorname{dom} h = \mathbb{Z}^G$.
2. **while** there exists $\varepsilon \in \{-1, 1\}$ and $S \subseteq G$ with $h(p + \varepsilon\mathbf{1}_S) < h(p)$ **do**
3. Choose S, ε such that $h(p + \varepsilon\mathbf{1}_S) - h(p)$ is minimized.
4. Set $p := p + \varepsilon\mathbf{1}_S$.
5. **end while**
6. p is a competitive equilibrium price vector.

Monotone Convergence

Although we have already seen that the algorithm converges globally, it is often desirable to have an algorithm that converges monotonically (i.e., for iterative combinatorial auctions). The following tie-breaking rule implies that if the algorithm starts with a price that is below every equilibrium price and always sets $\varepsilon = +1$, then it monotonically converges to the lowest equilibrium price:

$$\text{Choose the (unique) minimal minimizer } S \text{ of } h(p + \mathbf{1}_S) - h(p). \tag{4}$$

A unique minimal minimizer exists because h is submodular.

Let p^* be the lowest equilibrium price (such a price exists since the set of competitive equilibrium prices is a lattice). Convergence follows if the modified algorithm never stops strictly below p^* and always stays below p^*. The former property is a consequence

[5] In fact, this argument implies the convergence of *any* descent algorithm.

of the integral convexity of h; the latter property follows from the tie-breaking rule and the submodularity of h.

Lemma 1. *If the algorithm stops at p, then $p \geq p^*$.*

Proof. Assume that the algorithm stops at p but $p < p^*$. Then there exists $\lambda \in (0, 1)$ such that $p' = \lambda p + (1 - \lambda)p^*$ has the property that $||p' - p||_\infty < 1$. The integral neighborhood $N(p')$ consists of vectors $p + \mathbf{1}_X$ for subsets $X \subseteq G$. Since the algorithm stopped at p, $h(p+\mathbf{1}_X) \geq h(p)$ for all these subsets. Hence, by the definition of the local convex extension, $\tilde{h}(p') \geq \tilde{h}(p)$. But this is a contradiction because $\tilde{h}(p') < \tilde{h}(p) = h(p)$ due to the convexity of $\tilde{h}$, which holds by integral convexity of h. □

Lemma 2. *If $p' = p + \mathbf{1}_S$ is chosen by the algorithm at $p \leq p^*$ then $p' \leq p^*$.*

Proof. Submodularity of h implies

$$h(p^*) + h(p') \geq h(p^* \vee p') + h(p^* \wedge p').$$

Since p^* is an equilibrium price vector, we have $h(p^*) \leq h(p^* \vee p')$, and hence $h(p') \geq h(p^* \wedge p')$. Now assume that there is a $j \in G$ with $p'_j > p^*_j$. Then $p^* \wedge p' = p + \mathbf{1}_T$ for $T \subseteq S \setminus \{j\}$, which contradicts the minimality of S prescribed by the tie-breaking rule (4). □

For instance, if free-disposal can be assumed, then every competitive equilibrium price is always non-negative and hence the monotone tâtonnement process can always be started with a price of $p = 0$ and converges to the lowest equilibrium price.

We get the following summarizing result:

Theorem 5. *If the valuation function of every agent $i \in N$ satisfies Assumption 1, then for every endowment $\bar{x} \in \mathbb{Z}^G$ and every initial price vector $p \in \mathbb{Z}^G$, Algorithm 1 converges to a competitive equilibrium price vector.*

Further, if the starting price vector is lower than every equilibrium price, then the monotonic version of Algorithm 1 converges to the lowest competitive equilibrium price.

Eliciting Descent Directions

In practice, for instance when using the tâtonnement process as an iterative auction, it is impractical to elicit the values of the indirect utility functions $U_i(p)$ and therefore impossible to evaluate $h(p + \varepsilon\mathbf{1}_S)$. However, when every agent reports his demand correspondence $D_i(p)$ at the current price p, it is possible to compute the difference $h(p+\varepsilon\mathbf{1}_S) - h(p)$ for every ε and S (Ausubel 2006). The intuitive reason for this is that for an $L^\natural$-convex function h, the difference between the function values at $p + \varepsilon\mathbf{1}_S$ and p can be constructed from the subdifferential of h at the point p. The following lemma demonstrates that this subdifferential corresponds to the set of excess supply vectors.

Lemma 3. *For any $p \in \mathbb{Z}^N$ we have $x \in D_N(p) \;\Leftrightarrow\; \bar{x} - x \in \partial h(p)$.*

Proof. This is a consequence of the conjugacy between v_N and U_N. First note that by the definition of h, we have $\bar{x} - x \in \partial h(p) \;\Leftrightarrow\; -x \in \partial U_N(p)$. Next, $x \in D_N(p)$ is equivalent to $U_N(p) = v_N(x) - \langle p, x\rangle$. Since $v_N(x) = \inf_q\{U_N(q) + \langle q, x\rangle\}$ by conjugacy, this is equivalent to

$$U_N(p) + \langle p, x\rangle \leq U_N(q) + \langle q, x\rangle \quad \forall q \in \mathbb{Z}^N.$$

This is in turn just the definition of $-x \in \partial U_N(p)$, which completes the proof. □

Since h is an $L^\natural$-convex function, the difference between $h(p + \varepsilon \mathbf{1}_S)$ and $h(p)$ can be computed via the support function of its subgradient $\partial h(p)$ at p, evaluated in the direction of $\varepsilon \mathbf{1}_S$:

Lemma 4. *Let g be an $L^\natural$-convex function, $p \in \operatorname{dom} g$ and $S \subseteq G$, $\varepsilon \in \{-1, 1\}$. Then*

$$g(p + \varepsilon \mathbf{1}_S) - g(p) = \max_{y \in \partial g(p)} \langle y, \varepsilon \mathbf{1}_S\rangle.$$

For a proof see, for example, Proposition 7.44 in Murota (2003). The difference $h(p + \varepsilon \mathbf{1}_S) - h(p)$ can now be computed as follows. By Lemma 3,

$$\max_{y \in \partial h(p)} \langle y, \varepsilon \mathbf{1}_S\rangle = \max_{x \in D_N(p)} \langle \bar{x} - x, \varepsilon \mathbf{1}_S\rangle$$

and therefore, we have established the following proposition:

Proposition 6. *Assume that x^* solves the optimization problem*

$$\min_{x \in D_N(p)} \langle x, \varepsilon \mathbf{1}_S\rangle. \tag{5}$$

Then, by Lemma 4,

$$h(p + \varepsilon \mathbf{1}_S) - h(p) = \langle \bar{x} - x^*, \varepsilon \mathbf{1}_S\rangle.$$

The optimization problem (5) can be decomposed by solving $\min_{x^i \in D_i(p)} \langle x^i, \varepsilon \mathbf{1}_S\rangle$ for every agent separately and then setting $x^* = \sum_{i \in N} x^i$, since the objective function is linear and $D_N(p)$ is the Minkowski sum of the sets $D_i(p)$.[6] Also, since the sets $D_i(p)$ are $M^\natural$-convex, the greedy algorithm provides a way to maximize a linear objective function over $D_i(p)$ efficiently (Dress and Wenzel 1990).

The following example illustrates how the descent direction can be derived from the demand correspondence for the case of one agent.

Example 1. *Assume that there is one agent with the valuation function v depicted in Figure 5a and that the economy is endowed with $\bar{x} = 2$ units of one good. The corresponding function $h(p) = 2p + U(p)$ is shown in Figure 5b. Let the price adjustment process start with $p = 3$. At this price, the agent will demand quantities of either* 0 *or* 1*. In accordance with Lemma 3, the subdifferential of h at $p = 3$ is $\{1, 2\}$. Slopes in*

[6] Ausubel (2006) gives a different proof of a version of Proposition 6 which makes use of the single-improvement property that v_N satisfies since it is $M^\natural$-concave. Instead, we use the $L^\natural$-convexity of U_N and h.

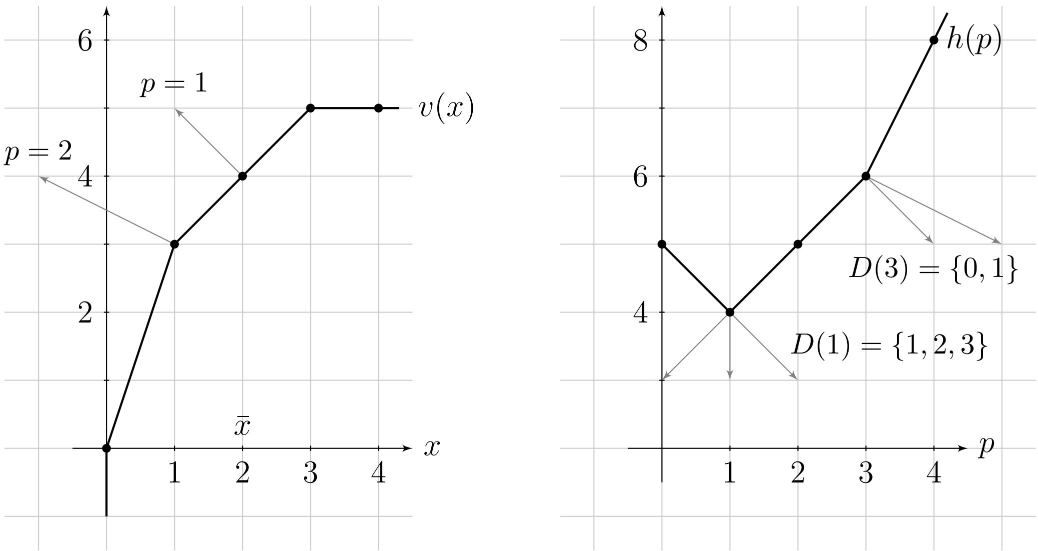

(a) Valuation function used in Example 1.

(b) The function h corresponding to the valuation function on the left.

Figure 5: Illustration of the relation between demand correspondence and subdifferential in Example 1.

the direction of -1 *and* 1 *are given by* $1 \cdot (-1) = -1$ *and* $2 \cdot 1 = 2$*, respectively, and therefore the price should be adjusted downwards. At* $p = 2$*, the agent only demands a quantity of* 1 *and the price should be lowered further. At* $p = 1$*, the agent demands* $D(1) = \{1, 2, 3\}$*, and the subdifferential of* h *at* $p = 1$ *is* $\{-1, 0, 1\}$*. Since the slope in the direction of* -1 *and* 1 *is* $(-1) \cdot (-1) = 1$ *and* $1 \cdot 1 = 1$*, respectively, we know that* $p = 1$ *is a minimum of the function* h *and that we have found an equilibrium.*

At this point we comment on the discrete price adjustment process which is given in Milgrom and Strulovici (2009). The authors state that their adjustment process, which is an approximation to a convergent continuous process described by an equation of motion in continuous time, only requires knowledge of one element of the subdifferential for any given price. Although the authors give results that there is always a price grid fine enough (or equivalently, a scaling of the valuation functions high enough) such that the trajectory of the discrete process lies in an ε-tube around the continuous trajectory, this does not mean that the discrete process converges. In fact, without knowledge of the complete demand sets at some price p, it is already impossible to tell whether p is a market-clearing price.[7]

While the method described in Proposition 6 can be used to evaluate $h(p + \varepsilon \mathbf{1}_S) - h(p)$, it remains to find $\varepsilon \in \{-1, 1\}$ and S such that this term is minimized. Since h is $L^\natural$-convex and in particular submodular, $h(p + \varepsilon \mathbf{1}_S) - h(p)$ is submodular in (ε, S), and efficient algorithms for minimizing submodular set functions can be used (Schrijver 2000).

[7] For instance, in Example 1, if the agent only reports to demand a quantity of 1 at $p = 1$, the auctioneer has to conclude that p is not an equilibrium price, although it in fact is.

Remark 2. *If one does not insist on a consecutive price trajectory but instead allows the ask price to jump around freely, the price adjustment algorithm can be scaled. The resulting algorithm then finds a competitive equilibrium in strongly polynomial time (see Murota 2003).*

7. Applications

In this section we demonstrate how different models from the literature fit in our framework and how our results can be applied to them.

Ausubel's Auction for Heterogeneous Goods

In a seminal paper, Ausubel (2006) developed a price adjustment process for auction settings with indivisible goods where agents' preferences satisfy the gross substitutes property. The auction proceeds through the minimization of a Lyapunov function and the analysis in this chapter is heavily inspired by this. Conversely, Ausubel's (2006) auction is an important special case of our adjustment process:

Corollary 1. *Assume that agents have valuation functions over the unit cube* $\{0,1\}^G$ *and that the economy is endowed with one unit of each good* ($\bar{x} = \mathbf{1}_G$)*. Then the tâtonnement process outlined in Section 6 describes the discrete price adjustment process presented in Ausubel (2006).*

Thus, our model generalizes Ausubel's model in that it works for preferences over arbitrary positive and/or negative quantities of every good, as well as any initial endowment. While positive quantities other than one can be simulated in Ausubel's framework by modeling every unit as a separate good, the auction then results in non-linear prices. In contrast, our algorithm generates linear prices for arbitrary quantities.

Milgrom and Strulovici (2009) also generalize Ausubel to multiple units of goods and introduce the strong substitutes condition, which, for positive quantities, is equivalent to Assumption 1. Therefore, our work also generalizes Milgrom and Strulovici (2009) by allowing for negative quantities and therefore for the possibility to model producers.

Our framework can also be applied to the set of preferences that are used in the Product-Mix Auction introduced by Klemperer (2010), since these preferences are a special case of gross substitutes.

Gross Substitutes and Complements

The double-track adjustment process presented in Sun and Yang (2009) is a special case of the price adjustment process outlined above. We start by recalling the gross substitutes and complements condition. In the model introduced by Sun and Yang (2006), the set of goods is partitioned into two sets $G = G_1 \sqcup G_2$.

Definition 9. *A valuation* $v_i : \{0,1\}^G \to \mathbb{R}$ *satisfies the* weak/ordinary gross substitutes and complements (GSC) *condition, if, given some price vector* $p \in \mathbb{R}^G$*, some good*[8]

[8] In this definition, a and b are set to 1 and 2, or 2 and 1, respectively.

$j \in G_a$, *and* $\delta > 0$, *the following holds: For every* $x \in D_i(p)$ *there exists* $x' \in D_i(p+\delta\mathbf{1}_j)$ *such that for all* $k \neq j$, *we have* $x_k \leq x'_k$ *if* $k \in G_a$ *and* $x_k \geq x'_k$ *if* $k \in G_b$.

We show that every valuation that satisfies the GSC condition can be transformed into a valuation that satisfies the GS condition by reversing the sign of every good in G_2. Assume that the goods are ordered such that the goods in G_1 come before the goods in G_2. Then the transformation can be described by applying the matrix

$$M = \begin{pmatrix} I_{|G_1|} & 0 \\ 0 & -I_{|G_2|} \end{pmatrix},$$

where $I_{|G_a|}$ is the identity matrix of dimension $|G_a|$. Using this transformation we can define the transformed valuation function M^*v_i through $M^*v_i(x) = v_i(Mx)$. The transformed indirect utility M^*U_i and demand correspondence M^*D_i are defined using the transformed valuation function.

Lemma 5. *We have* $x \in D_i(p)$ *if and only if* $M^{-1}x \in M^*D_i(Mp)$.

Proof. By definition, we have $x \in D_i(p)$ if and only if $v_i(x) - \langle p, x\rangle \geq v_i(x') - \langle p, x'\rangle$ for all $x' \in \operatorname{dom} v_i$. By substituting $x = My$ and $x' = My'$, this is equivalent to

$$\begin{aligned} & v_i(My) - \langle p, My\rangle \geq v_i(My') - \langle p, My'\rangle \\ \Leftrightarrow \quad & M^*v_i(y) - \langle Mp, y\rangle \geq M^*v_i(y') - \langle Mp, y'\rangle \quad \forall y \in \operatorname{dom} M^*v_i, \end{aligned}$$

which in turn means that $y = M^{-1}x \in M^*D_i(Mp)$. □

Proposition 7. *Let* $v_i : \{0,1\}^G \to \mathbb{R}$. *Then* v_i *satisfies GSC if and only if* M^*v_i *satisfies wGS (i.e., is* $M^\natural$*-concave).*

Proof. First note that wGS is equivalent to a version where the price of only one good is increased. We show equivalence to this modified definition.

Assume that v_i satisfies GSC. Let $p \in \mathbb{R}^G$, $\delta > 0$ and $j \in G_1$. Define $p' = p + \delta\mathbf{1}_j$ and let $x \in M^*D_i(p)$. We need to find $x' \in M^*D_i(p')$ such that for $k \neq j$, $x_k \leq x'_k$. By Lemma 5 we know that $y = Mx \in D_i(Mp)$. Also, since $j \in G_1$, $M(p+\delta\mathbf{1}_j) = Mp+\delta\mathbf{1}_j$. Since v_i satisfies the GSC condition, we know that there exists $y' \in D_i(Mp + \delta\mathbf{1}_j)$ such that for all $k \neq j$, we have $y_k \leq y'_k$ if $k \in G_1$ and $y_k \geq y'_k$ if $k \in G_2$.

We claim that $x' = M^{-1}y'$ satisfies the requirements. First, by Lemma 5, $x' \in M^*D_i(p + \delta\mathbf{1}_j)$. Now take some good $k \neq j$. If $k \in G_1$ then $x_k = y_k$ and $x'_k = y'_k$ and therefore $x_k \leq x'_k$. If $k \in G_2$, then $x_k = -y_k$ and $x'_k = -y'_k$ and therefore $x_k = -y_k \leq -y'_k = x'_k$.

The argument is similar for the case where $j \in G_2$ and also sufficiency can be shown analogously. □

Proposition 7 motivates the following definition of generalized gross substitutes and complements for multiple units of goods (cf. Baldwin and Klemperer 2013):

Definition 10. *Valuation* v_i *satisfies the* generalized gross substitutes and complements (GGSC) *condition, if* M^*v_i *is* $M^\natural$*-concave.*

With this definition, existence of competitive equilibria follows immediately: For a set of valuation functions $\{v_i\}_{i \in N}$ that satisfy GGSC and endowment $\bar{x}$, consider the modified economy $\{M^* v_i\}_{i \in N}$ with endowment $M\bar{x}$. Since $\{M^* v_i\}_{i \in N}$ are $M^\natural$-concave, there exists a competitive equilibrium price vector p^*, that is, $M\bar{x} \in M^* D_N(p^*)$. Then, by Lemma 5, $\bar{x} \in D_N(Mp^*)$, so Mp^* is a competitive equilibrium price vector for the original economy.

We note that the application of our results via the described transformation above requires us to be able to deal with non free disposal valuations. Specifically, if a valuation v_i satisfies free disposal then $M^* v_i$ has "anti free disposal" for goods in G_2. It follows as in Proposition 3 that the price p_j^* for $j \in G_2$ is non-positive, and therefore Mp_j^* is non-negative.

The above transformation also allows us to describe the double-track price adjustment process by Sun and Yang (2006) in terms of the algorithm from Section 6: The algorithm is run on the modified economy $\{M^* v_i\}_{i \in N}$ and $M\bar{x}$ (call it *internal representation*). If we transform this algorithm back to the original economy (call it *external representation*), we get the price adjustment process described in Sun and Yang (2009). In particular,

(i) if the current internal price is Mp, the price p is presented to the agents. If the internal price for some good in G_2 increases, then the external price decreases and vice versa.

(ii) if an agent indicates that he demands bundle x, then $M^{-1}x$ is used for the internal calculation of the next price.

(iii) if the monotone convergence algorithm is used internally, the starting price has to be set such that it is below every competitive equilibrium price. In the original economy, this means that the price for goods in G_1 has to be set to the lowest and the price for goods in G_2 to the highest possible level. Then, since the algorithm converges monotonically in the internal representation, this means that the real price for goods in G_1 increases whereas the real price for goods in G_2 decreases.

(iv) since the set of (internal) equilibrium prices is a lattice, the set of transformed equilibrium prices forms a "generalized lattice" as defined by Sun and Yang (2009).

Hence, we can formulate the following corollary of Theorems 4 and 5.

Corollary 2. *Assume that agents have valuation functions over the unit cube* $\{0,1\}^G$ *and that these valuation functions satisfy GSC. Further, assume that* $\bar{x} = \mathbf{1}_G$*. Then, a competitive equilibrium exists and the procedure outlined above describes the double-track adjustment process presented in Sun and Yang (2009).*

Thus, the results in this chapter generalize Sun and Yang (2006) as well as Sun and Yang (2009) in that they work for preferences over arbitrary positive and/or negative quantities of every good, as well as for any initial endowment of the economy, if the valuation functions satisfy the GGSC condition.

Trading Networks

The trading networks economy introduced by Hatfield et al. (2013) also fits into our model. We first describe the network economy and then show how the valuation functions in this chapter relate to the valuation functions as they are defined in Hatfield et al. (2013).

In the model, there is a set of agents N and a set of trades Ω which can be interpreted as goods. The agents and trades form a graph, where the nodes are the agents and each trade is a directed edge. If trade $\omega = (i, i') \in \Omega$ points from agent i to agent i' then we say that agent i is the seller and agent i' is the buyer in this trade. Let Ω_i be the trades that are adjacent to agent i. Every agent i has a valuation function $v_i : \{-1, 0, 1\}^{\Omega_i} \to \mathbb{R}$ over subsets of adjacent trades. We model agent i being a buyer in trade $\omega \in \Omega_i$ by requiring that for $x \in \{-1, 0, 1\}^{\Omega_i}$, $v_i(x) = -\infty$ if $x_\omega = -1$. Similarly, if agent i is a seller in trade ω we require $v_i(x) = -\infty$ whenever $x_\omega = 1$. The interpretation is that if agent i demands vector x with $x_\omega = -1$ then he wants to be engaged in trade ω where he is the seller and similarly, if $x_\omega = 1$ then he wants to be engaged in trade ω where he is the buyer.

We can embed this economy in our model by extending the valuation functions v_i to $\{-1, 0, 1\}^{\Omega}$ as follows: Set $v_i(x) = -\infty$ if $x_\omega \neq 0$ for some $\omega \notin \Omega_i$. Otherwise, if $x_\omega = 0$ for all $\omega \notin \Omega_i$, just copy the valuation of the vector x restricted to Ω_i. A competitive equilibrium in the network economy is a competitive equilibrium for the endowment $\bar{x} = 0$. Then we know that, whenever all valuation functions v_i satisfy Assumption 1, there exists a competitive equilibrium and a convergent price adjustment process.

We therefore get the following corollary regarding the model by Hatfield et al. (2013):

Corollary 3. *In the model defined above, if the valuation function of every agent satisfies Assumption 1, a competitive equilibrium exists. Further, Algorithm 1 can be used to find competitive equilibrium prices for any initial price vector p.*

In the following we explain how Assumption 1 is equivalent to the full substitutes condition defined in Hatfield et al. (2013). In their paper, valuation functions, utility functions, demand, and indirect utility are defined slightly differently as follows: Every agent i has a valuation function $\tilde{v}_i : \{0, 1\}^{\Omega_i} \to \mathbb{R}$ that can be embedded into $\{0, 1\}^{\Omega}$ as described above. Let $\Omega_{i\to}$ be the trades adjacent to agent i in which he is a seller and let $\Omega_{\to i}$ be the trades adjacent to him in which he is a buyer, respectively. Then the interpretation is that if agent i demands bundle x and $x_\omega = 1$ then agent i wants to be engaged in trade ω. Given price vector $p \in \mathbb{R}^{\Omega}$, an agent's quasi-linear utility is defined as

$$\tilde{u}_i(x, p) = \tilde{v}_i(x) + \sum_{\omega \in \Omega_{i\to} : x_\omega = 1} p_\omega - \sum_{\omega \in \Omega_{\to i} : x_\omega = 1} p_\omega.$$

Indirect utility $\tilde{U}_i$ and demand correspondence $\tilde{D}_i$ are then defined as in Section 2, but using $\tilde{u}_i$.

Hatfield et al. (2013) assume full substitutability which is defined as follows:[9]

Definition 11. *A valuation function v_i satisfies* full substitutability (FS) *if for every two price vectors $p \leq p'$ the following holds: For every $x \in \tilde{D}_i(p)$ there exists $x' \in \tilde{D}_i(p')$*

[9]This formulation is similar and equivalent to "indicator language full substitutability."

such that whenever $p_\omega = p'_\omega$ for some ω, then $x_\omega \leq x'_\omega$ if $\omega \in \Omega_{\to i}$ and $x_\omega \geq x'_\omega$ if $\omega \in \Omega_{i\to}$.

We now show that full substitutability and gross substitutability are equivalent. Fix some agent i. We introduce the following transformation of a vector $x \in \{-1, 0, 1\}^\Omega$. As in the last subsection, assume that the trades are ordered such that trades $\omega \in \Omega_{i\to}$ come first. Then we apply the following matrix:

$$M = \begin{pmatrix} -I_{|\Omega_{i\to}|} & 0 \\ 0 & I_{|\Omega\setminus\Omega_{i\to}|} \end{pmatrix}$$

From the interpretation of the valuation functions v_i and $\tilde{v}_i$ we see that for them to represent the same preferences over trades, $v_i(Mx) = \tilde{v}_i(x)$ has to hold for all $x \in \{0, 1\}^\Omega$.

Lemma 6. *For transformed bundles we have $Mx \in D_i(p) \Leftrightarrow x \in \tilde{D}_i(p)$.*

Proof. This follows from

$$\sum_{\omega\in\Omega_{i\to}:x_\omega=1} p_\omega - \sum_{\omega\in\Omega_{\to i}:x_\omega=1} p_\omega = -\langle p, Mx\rangle$$

and $v_i(Mx) = \tilde{v}_i(x)$. □

Proposition 8. *A valuation $\tilde{v}_i$ satisfies FS if and only if the corresponding valuation function v_i satisfies Assumption 1.*

Proof. On the unit-cube, Assumption 1 is equivalent to the ordinary (weak) gross substitutes condition. After applying the translation $v'_i(x) = v_i(x - \mathbf{1}_{\Omega_{i\to}})$, $\operatorname{dom} v'_i$ is the unit-cube. Since the gross substitutes condition is translation-invariant, it is therefore enough to show that $\tilde{v}_i$ satisfies FS if and only if v_i satisfies weak GS (Definition 2).

In order to prove necessity assume that $\tilde{v}_i$ satisfies FS. Take price vectors $p \leq p'$ and $x \in D_i(p)$. By Lemma 6 we know that $Mx \in \tilde{D}_i(p)$. Since $\tilde{v}_i$ satisfies FS, we know that there exists $Mx' \in \tilde{D}_i(p')$ such that whenever $p_\omega = p'_\omega$ for some ω, then $Mx_\omega \leq Mx'_\omega$ if $\omega \in \Omega_{\to i}$ and $Mx_\omega \geq Mx'_\omega$ if $\omega \in \Omega_{i\to}$. Hence, for ω with $p_\omega = p'_\omega$ we have $x_\omega \leq x'_\omega$. Furthermore, $x' \in D_i(p')$ and therefore x' satisfies all requirements in Definition 2.

Sufficiency is proved similarly. □

Combining the transformation M with the translation v' in the proof above yields the same transformation as that which is used in the proof for the existence of competitive equilibria in Hatfield et al. (2013). However, the translation is not needed in our model since our framework can deal with negative amounts of goods (that is, producers).

We can also use our framework to extend the model in this subsection to multiple units of goods in each trade. Further, the transformation in the last subsection on the gross substitutes and complements condition can be applied to the trading network model to get two sets of trades Ω^1 and Ω^2 where trades in the same sets substitute each other but trades in different sets are complements (see also Drexl 2013).

8. Discussion

We have used the theory of Discrete Convex Analysis to unify and generalize the literature on tâtonnement for economies with indivisibilities. The interpretation of the auction procedure proposed by Ausubel (2006) as a steepest descent algorithm of certain discrete convex functions yields simple and intuitive proofs of the convergence properties of the generalized adjustment process. Applying the results demonstrates that all the literature on discrete tâtonnement harnesses the notion of gross substitutability, which is equivalent to $M^\natural$-convexity. The theory of Discrete Convex Analysis confirms that $M^\natural$-convexity is essential for many properties of discrete convex functions.

Since the existence of market-clearing equilibria can be guaranteed for classes of valuation functions that are much more general than the valuations we consider (see Baldwin and Klemperer 2013), one of the big open questions is whether a price adjustment process can be designed that converges for every instance of valuation functions where equilibria are guaranteed to exist. While the indirect utility function in these cases is still convex, it does not exhibit all the combinatorial properties used in the present chapter. Still, it might very well be possible to prove convergence of a suitably designed descent or gradient algorithm that minimizes the aggregate indirect utility function.

References

Arrow, K. (1979). The property rights doctrine and demand revelation under incomplete information, *in* M. Boskin (ed.), *Economics and Human Welfare*, Academic Press, New York.

Arrow, K. J. and Debreu, G. (1954). Existence of equilibrium for a competitive economy, *Econometrica* **22**: 265–290.

Arrow, K. J. and Hurwicz, L. (1958). On the stability of the competitive equilibrium, I, *Econometrica* **26**: 522–552.

Athey, S. and Miller, D. A. (2007). Efficiency in repeated trade with hidden valuations, *Theoretical Economics* **2**(3): 299–354.

Ausubel, L. M. (2006). An efficient dynamic auction for heterogeneous commodities, *American Economic Review* **96**: 602–629.

Azrieli, Y. and Kim, S. (2012). Pareto efficiency and weighted majority rules. Working Paper, Ohio State University.

Baldwin, E. and Klemperer, P. (2013). Tropical geometry to analyse demand. Working Paper, Oxford University.

Barbera, S. and Jackson, M. O. (2004). Choosing how to choose: Self-stable majority rules and constitutions, *Quarterly Journal of Economics* **119**(3): 1011–1048.

Barbera, S. and Jackson, M. O. (2006). On the weights of nations: Assigning voting weights in a heterogeneous union, *Journal of Political Economy* **114**(2): 317–339.

Bergemann, D. and Morris, S. (2005). Robust mechanism design, *Econometrica* **73**(6): 1771–1813.

Bergemann, D. and Välimäki, J. (2010). The dynamic pivot mechanism, *Econometrica* **78**(2): 771–789.

Bierbrauer, F. and Hellwig, M. (2012). Mechanism design and voting for public-good provision, *Technical report*.

Buchanan, J. M. and Tullock, G. (1962). *The calculus of consent: logical foundations of constitutional democracy*, University of Michigan Press.

Carrasco, V. and Fuchs, W. (2011). From equals to despots: The dynamics of repeated decision taking in partnerships with private information, *Technical report*.

Casella, A. (2005). Storable votes, *Games and Economic Behavior* **51**(2): 391–419.

Casella, A., Llorente-Saguer, A. and Palfrey, T. R. (2012). Competitive equilibrium in markets for vots, *Journal of Political Economy* **120**(4): 593–658.

Cavallo, R. (2006). Optimal decision-making with minimal waste: Strategyproof redistribution of vcg payments, *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '06, ACM, New York, NY, USA, pp. 882–889.

Chakravarty, S. and Kaplan, T. R. (2013). Optimal allocation without transfer payments, *Games and Economic Behavior* **77**(1): 1–20.

Chawla, S., Hartline, J. D., Rajan, U. and Ravi, R. (2006). Bayesian optimal no-deficit mechanism design, *Internet and Network Economics* **4286**: 136–148.

Condorelli, D. (2012). What money can't buy: Efficient mechanism design with costly signals, *Games and Economic Behavior* **75**(2): 613–624.

Danilov, V., Koshevoy, G. and Lang, C. (2003). Gross substitution, discrete convexity and submodularity, *Discrete Applied Mathematics* **131**: 283–298.

Danilov, V., Koshevoy, G. and Murota, K. (2001). Discrete convexity and equilibria in economies with indivisible goods and money, *Mathematical Social Sciences* **41**: 251–273.

d'Aspremont, C. and Gerard-Varet, L.-A. (1979). Incentives and incomplete information, *Journal of Public Economics* **11**(1): 25–45.

de Clippel, G., Naroditskiy, V., Polukarov, M., Greenwald, A. and Jennings, N. (2013). Destroy to save, *Games and Economic Behavior* . Forthcoming.

Demange, G., Gale, D. and Sotomayor, M. (1986). Multi-item auctions, *Journal of Political Economy* **94**: 863–872.

deVries, S., Schummer, J. and Vohra, R. V. (2007). On ascending vickrey auctions for heterogeneous objects, *Journal of Economic Theory* **132**(1): 95–118.

Diakonikolas, I., Papadimitriou, C., Pierrakos, G. and Singer, Y. (2012). Efficiency-revenue trade-offs in auctions, *Automata, Languages, and Programming*, Vol. 7392 of *Lecture Notes in Computer Science*, pp. 488–499.

Dress, A. W. M. and Wenzel, W. (1990). Valuated matroid: A new look at the greedy algorithm, *Applied Mathematics Letters* **3**: 33–35.

Drexl, M. (2013). Substitutes and complements in trading networks. Working Paper, University of Bonn.

Drexl, M. and Kleiner, A. (2012). Optimal private good allocation: The case for a balanced budget. Discussion Paper, University of Bonn.

Esö, P. and Futo, G. (1999). Auction design with a risk averse seller, *Economics Letters* **65**(1): 71–74.

Fujishige, S. and Yang, Z. (2003). A note on Kelso and Crawford's gross substitutes condition, *Mathematics of Operations Research* **28**: 463–469.

Gale, D. (1960). *The Theory of Linear Economic Models*, University of Chicago Press.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage, *The American Mathematical Monthly* **69**(1): 9–15.

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

Gershkov, A., Moldovanu, B. and Shi, X. (2013). Optimal mechanism design without money. Working Paper, University of Bonn.

Gibbard, A. (1973). Manipulation of voting schemes: A general result, *Econometrica* **41**(4): 587–601.

Green, J. and Laffont, J.-J. (1977). Révélation des préférences pour les biens public: Caractérisation des mécanismes satisfaisants, *Cahiers du Séminaire d'Économétrie* **1**(19): 83–103.

Green, J. R. and Laffont, J.-J. (1979). *Incentives in Public Decision-Making*, North-Holland Pub. Co., Amsterdam, NL.

Groves, T. (1973). Incentives in teams, *Econometrica* **41**(4): 617–631.

Groves, T. and Loeb, M. (1975). Incentives and public inputs, *Journal of Public Economics* **4**(3): 211–226.

Gul, F. and Stacchetti, E. (1999). Walrasian equilibrium with gross substitutes, *Journal of Economic Theory* **87**(1): 95–124.

Guo, M. and Conitzer, V. (2008). Better redistribution with inefficient allocation in multi-unit auctions with unit demand, *Proceedings of the 9th ACM conference on Electronic commerce*, EC '08, ACM, pp. 210–219.

Guo, M. and Conitzer, V. (2009). Worst-case optimal redistribution of vcg payments in multi-unit auctions, *Games and Economic Behavior* **67**(1): 69–98.

Guo, M. and Conitzer, V. (2010). Optimal-in-expectation redistribution mechanisms, *Artificial Intelligence* **174**(5–6): 363–381.

Hagerty, K. M. and Rogerson, W. P. (1987). Robust trading mechanisms, *Journal of Economic Theory* **42**(1): 94–107.

Hartline, J. D. and Roughgarden, T. (2008). Optimal mechanism design and money burning, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, STOC '08, ACM, pp. 75–84.

Hatfield, J. W., Kominers, S. D., Nichifor, A., Ostrovsky, M. and Westkamp, A. (2013). Stability and competitive equilibrium in trading networks, *Journal of Political Economy* . Forthcoming.

Hortala-Vallve, R. (2010). Inefficiencies on linking decisions, *Social Coice and Welfare* **34**: 471–486.

Hortala-Vallve, R. (2012). Qualitative voting, *Journal of Theoretical Politics* **24**: 526–554.

Jackson, M. O. and Sonnenschein, H. F. (2007). Overcoming incentive constraints by linking decisions, *Econometrica* **75**(1): 241–257.

Kelso, A. S. and Crawford, V. P. (1982). Job matching, coalition formation, and gross substitutes, *Econometrica* **50**(6): 1483–1504.

Klemperer, P. (2010). The product-mix auction: A new auction design for differentiated goods, *Journal of the European Economic Association* **8**: 526–536.

Koopmans, T. and Beckmann, M. (1957). Assignment problems and the location of economic activities, *Econometrica* **25**: 53–76.

Laffont, J.-J. and Maskin, E. (1982). Advances in economic theory, *in* W. Hildenbrand (ed.), *Advances in Economic Theory*, Cambridge University Press, Cambridge, UK, pp. 31–94.

Ledyard, J. O. and Palfrey, T. R. (2002). The approximation of efficient public good mechanisms by simple voting schemes, *Journal of Public Eocnomics* **83**(2): 153–171.

Manelli, A. M. and Vincent, D. R. (2010). Bayesian and dominant-strategy implementation in the independent private-values model, *Econometrica* **78**(6): 1905–1938.

McAfee, R. P. (1992). A dominant strategy double auction, *Journal of Economic Theory* **56**(2): 434–450.

McAfee, R. P. and McMillan, J. (1992). Bidding rings, *The American Economic Review* **82**(3): 579–599.

McKenzie, L. (1959). On the existence of general equilibrium for a competitive market, *Econometrica* **27**: 54–71.

Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction, *Journal of Political Economy* **108**: 245–272.

Milgrom, P. and Strulovici, B. (2009). Substitute goods, auctions, and equilibrium, *Journal of Economic Theory* **144**: 212–247.

Miller, D. (2012). Robust collusion with private information, *Review of Economic Studies* **79**(2): 778–811.

Moulin, H. (1983). *The Strategy of Social Choice*, North-Holland Pub. Co., Amsterdam, NL.

Moulin, H. (2009). Almost budget-balanced vcg mechanisms to assign multiple objects, *Journal of Economic Theory* **144**(1): 96–119.

Murota, K. (2003). *Discrete Convex Analysis*, SIAM Society for Industrial and Applied Mathematics, Philadelphia.

Murota, K. and Tamura, A. (2002). New characterizations of M-convex functions and their applications to economic equilibrium models with indivisibilities, *Discrete Applied Mathematics* **131**: 495–512.

Myerson, R. B. (1981). Optimal auction design, *Mathematics of Operations Research* **6**(1): 58–73.

Myerson, R. B. (1986). Multistage games with communication, *Econometrica* **54**(2): 323–358.

Myerson, R. B. and Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading, *Journal of Economic Theory* **29**(2): 265–281.

Nehring, K. (2004). The veil of public ignorance, *Journal of Economic Theory* **119**(2): 247–270.

Ostrovski, M. (2008). Stability in supply chain networks, *American Economic Review* **98**(3): 897–923.

Parkes, D. C. and Ungar, L. H. (2000). Iterative combinatorial auctions: Theory and practice, *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pp. 74–81.

Pavan, A., Segal, I. and Toikka, J. (2008). Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure, *Carlo Alberto Notebooks 84*, Collegio Carlo Alberto.

Rae, D. W. (1969). Decision-rules and individual values in constitutional choice, *American Political Science Review* **63**(1): 40–56.

Riker, W. H. and Brams, S. J. (1973). The paradox of vote trading, *American Political Science Review* **67**(4): 1235–1247.

Samuelson, P. A. (1947). *Foundations of Economic Analysis*, Harvard University Press, Cambridge, MA.

Schmitz, P. and Tröger, T. (2012). The (sub-)optimality of the majority rule, *Games and Economic Behavior* **74**(2): 651–665.

Schrijver, A. (1986). *Theory of Linear and Integer Programming*, Wiley-Interscience, New York, NY.

Schrijver, A. (2000). A combinatorial algorithm minimizing submodular functions in strongly polynomial time, *Journal of Combinatorial Theory* **80**: 346–355.

Schwartz, J. A. and Wen, Q. (2012). Robust trading mechanisms with budget surplus and partial trade. Working Paper, Kennesaw State University and Vanderbilt University.

Shao, R. and Zhou, L. (2012). Optimal allocation of an indivisible good. Working Paper, Yeshiva University and Shanghai Jiao Tong University.

Shapley, L. S. and Shubik, M. (1971). The assignment game I: The core, *International Journal of Game Theory* **1**: 111–130.

Sun, N. and Yang, Z. (2006). Equilibria and indivisibilities: Gross substitutes and complements, *Econometrica* **74**(5): 1385–1402.

Sun, N. and Yang, Z. (2009). A double-track adjustment process for discrete markets with substitutes and complements, *Econometrica* **77**(3): 933–952.

Tatur, T. (2005). On the trade off between deficit and inefficiency and the double auction with a fixed transaction fee, *Econometrica* **73**(2): 517–570.

Tideman, T. N. and Tullock, G. (1976). A new and superior process for making social choices, *Journal of Political Economics* **84**(6): 1145–1159.

Walker, M. (1980). On the nonexistence of a dominant strategy mechanism for making optimal public decisions, *Econometrica* **48**(6): 1521–1540.

Walras, L. (1874). *Elements of Pure Economics*, Irwin, Homewood, IL.

Yoon, K. (2011). Optimal mechanism design when both allocative efficiency and expenditure inefficiency matter, *Journal of Mathematical Economics* **47**(6): 670–676.