Characterization of Revenue Equivalence

Birgit Heydenreich[1] Rudolf Müller[1] Marc Uetz[2] Rakesh Vohra[3]

July 1, 2008

Abstract

The property of an allocation rule to be implementable in dominant strategies by a unique payment scheme is called *revenue equivalence*. We give a characterization of revenue equivalence based on a graph theoretic interpretation of the incentive compatibility constraints. The characterization holds for any (possibly infinite) outcome space and many of the known results are immediate consequences. Moreover, revenue equivalence can be identified in cases where existing theorems are silent.

1 Introduction

One of the most important results of auction theory is the Revenue Equivalence Theorem. Subject to certain reasonable assumptions, it concludes that a variety of different auctions generate the same expected revenue for the seller. Klemperer (1999) writes that "much of auction theory can be understood in terms of this theorem.....". Hence the long line of papers that have attempted to relax the sufficient conditions under which revenue equivalence holds. The present paper provides necessary and sufficient conditions on the underlying primitives for revenue equivalence to hold.

We consider direct revelation mechanisms for agents with multidimensional types. Such mechanisms consist of an allocation rule and a payment scheme. The allocation rule selects an outcome depending on the agents' reported types, whereas the payment scheme assigns a payment to every agent. We focus attention on allocation rules that are implementable in dominant strategies.[4] Call such rules implementable. In this environment we characterize the

[1]Maastricht University, Quantitative Economics, P.O.Box 616, 6200 MD Maastricht, The Netherlands. Email:{b.heydenreich,r.muller}@ke.unimaas.nl
The first author is supported by NWO grant 2004/03545/MaGW 'Local Decisions in Decentralised Planning Environments'.
[2]University of Twente, Applied Mathematics, P.O. Box 217, 7500 AE Enschede, The Netherlands. Email: m.uetz@utwente.nl

[3]Northwestern University, Department of Managerial Economics and Decision Sciences, Kellogg Graduate School of Management, Evanston IL 60208. Email: r-vohra@kellogg.northwestern.edu

[4]With appropriate adjustments our characterization of revenue equivalence holds for ex-post as well as Bayes-Nash incentive compatibility.

uniqueness of the relevant payment scheme in terms of conditions that are easily verified in potential applications. The property of an allocation rule to be implementable in dominant strategies by a unique payment scheme is called *revenue equivalence*. Our characterization of revenue equivalence is based on a graph theoretic interpretation of the incentive compatibility constraints. This interpretation has been used before by Rochet (1987), Gui, Müller, and Vohra (2004), Saks and Yu (2005) and Müller, Perea, and Wolf (2007) to identify allocation rules that are implementable in dominant strategies or in Bayes-Nash. The characterization holds for any (possibly infinite) outcome space.

The bulk of prior work on revenue equivalence (Green and Laffont (1977), Holmström (1979), Myerson (1981), Krishna and Maenner (2001), Milgrom and Segal (2002)) identifies sufficient conditions on the type space that ensure all allocation rules from a given class satisfy revenue equivalence. We know of only two papers that identify necessary as well as sufficient conditions—i.e. *characterizing* conditions—for revenue equivalence to hold. If the outcome space is finite, Suijs (1996) characterizes type spaces and valuation functions for which utilitarian maximizers satisfy revenue equivalence. Chung and Olszewski (2007) characterize type spaces and valuation functions for which *every* implementable allocation rule satisfies revenue equivalence, again under the assumption of a finite outcome space. Our characterization identifies for general outcome spaces a joint condition on the type space, valuation function *and* implementable allocation rule that characterize revenue equivalence. Given agents' type spaces and valuation functions, several allocation rules may be implementable in dominant strategies, some of which satisfy revenue equivalence and some do not. In such cases, all previous results have no bite. However, our characterization can be used to determine which of the allocation rules do satisfy revenue equivalence.

The remainder of the paper is organized as follows. In Section 2 we introduce notation and basic definitions. In Section 3, we derive our graph-theoretic characterization of revenue equivalence. In Section 4 we briefly discuss how our characterization applies in various settings.

2 Setting and Basic Concepts

Denote by $\{1,\ldots,n\}$ the set of *agents* and let A be the set of possible *outcomes*. Outcome space A is allowed to have infinitely many, even uncountably many, elements. Denote the *type* of agent $i \in \{1,\ldots,n\}$ by t_i. Let T_i be the *type space* of agent i. Type spaces T_i can be arbitrary sets. Agent i's preferences over outcomes are modeled by the *valuation function* $v_i\colon A \times T_i \to \mathbb{R}$, where $v_i(a,t_i)$ is the valuation of agent i for outcome a when he has type t_i.

A *mechanism* (f,π) consists of an *allocation rule* $f\colon \Pi_{i=1}^n T_i \to A$ and a *payment scheme* $\pi\colon \Pi_{i=1}^n T_i \to \mathbb{R}^n$. In a *direct revelation mechanism*, the allocation rule chooses for a vector t of aggregate type reports of all agents an outcome $f(t)$, whereas the payment scheme assigns a payment $\pi_i(t)$ to each agent i. Let the vector (t_i, t_{-i}) denote the aggregate type report vector when i reports t_i and the other agents' reports are represented by t_{-i}. We assume *quasi-linear utilities*, that is, the utility of agent i when the aggregate report vector is (t_i, t_{-i}) is $v_i(f(t_i, t_{-i}), t_i) - \pi_i(t_i, t_{-i})$.

Definition 1 (dominant strategy incentive compatible) *A direct revelation mechanism* (f, π) *is called* dominant strategy incentive compatible *if for every agent* i*, every type* $t_i \in T_i$*, all aggregate type vectors* t_{-i} *that the other agents could report and every type* $s_i \in T_i$ *that* i *could report instead of* t_i*:*

$$v_i(f(t_i, t_{-i}), t_i) - \pi_i(t_i, t_{-i}) \geq v_i(f(s_i, t_{-i}), t_i) - \pi_i(s_i, t_{-i}).$$

If for allocation rule f *there exists a payment scheme* π *such that* (f, π) *is a dominant strategy incentive compatible mechanism, then* f *is called* implementable in dominant strategies*, in short* implementable*.*

In this paper we assume that the allocation rule is implementable in dominant strategies and study the uniqueness of the corresponding payment scheme. We refer to the latter as revenue equivalence.[5]

Definition 2 (Revenue Equivalence) *An allocation rule* f *implementable in dominant strategies satisfies the* revenue equivalence *property if for any two dominant strategy incentive compatible mechanisms* (f, π) *and* (f, π') *and any agent* i *there exists a function* h_i *that only depends on the reported types of the other agents* t_{-i} *such that*

$$\forall t_i \in T_i : \pi_i(t_i, t_{-i}) = \pi'_i(t_i, t_{-i}) + h_i(t_{-i}).$$

3 Characterization of Revenue Equivalence

We give a necessary and sufficient condition for revenue equivalence with the aid of a graph theoretic interpretation used by Rochet (1987), Gui, Müller, and Vohra (2004) and Saks and Yu (2005) to characterize implementable allocation rules. We also adopt some of their notation.

Fix agent i and the reports, t_{-i}, of the other agents. For simplicity of notation we write T and v instead of T_i and v_i. Similarly, for any mechanism (f, π), we regard f and π as functions of i's type alone, i.e. $f : T \to A$ and $\pi : T \to \mathbb{R}$. If (f, π) is dominant strategy incentive compatible, it is easy to see that for any pair of types $s, t \in T$ such that $f(t) = f(s) = a$ for some $a \in A$, the payments must be equal, i.e. $\pi(t) = \pi(s) =: \pi_a$. Hence, the payment of agent i is completely defined if the numbers π_a are defined for all outcomes $a \in A$ such that $f^{-1}(a)$ is nonempty. For

[5] We choose the term revenue equivalence in accordance with Krishna (2002). In our setting it is equivalent to payoff equivalence as used in Krishna and Maenner (2001). See Milgrom (2004), Section 4.3.1. for settings where it is not equivalent.

ease of notation, we let A denote the set of outcomes that can be a achieved for some report of agent i. For an allocation rule f, let us define two different kinds of graphs. The *type graph* T_f has node set T and contains an arc from any node s to any other node t of length[6]

$$\ell_{st} = v(f(t), t) - v(f(s), t).$$

Here, ℓ_{st} represents the gain in valuation for an agent truthfully reporting type t instead of lying type s. This could be positive or negative. The *allocation graph* G_f has node set A. Between any two nodes $a, b \in A$, there is a directed arc with length[6]

$$\ell_{ab} = \inf_{t \in f^{-1}(b)} (v(b, t) - v(a, t)).$$

The arc lengths ℓ_{ab} in the allocation graph represent the least gain in valuation for an agent with *any* type $t \in f^{-1}(b)$ for reporting truthfully, instead of misreporting so as to get outcome a (instead of b). The type graph and allocation graph are complete, directed, and possibly infinite graphs[7]. We introduce our main results in terms of allocation graphs. Analogous results hold for type graphs as well.

A *path* from node a to node b in G_f, or (a, b)-path for short, is defined as $P = (a = a_0, a_1, \dots, a_k = b)$ such that $a_i \in A$ for $i = 0, \dots, k$. Denote by *length*(P) the length of this path. A *cycle* is a path with $a = b$. For any a, we regard the path from a to a without any arcs as an (a, a)-path and define its length to be 0. Define $\mathcal{P}(a, b)$ to be the set of all (a, b)-paths.

Definition 3 (Node Potential) *A node potential p is a function $p\colon A \to \mathbb{R}$ such that for all $x, y \in A$, $p(y) \leq p(x) + \ell_{xy}$.*

Observation 1 *Let f be an allocation rule. Payment schemes π such that (f, π) is a dominant strategy incentive compatible mechanism, exactly correspond to node potentials in each of the allocation graphs G_f that are obtained from a combination of an agent and a report vector of the other agents.*

Proof. Assume f is implementable. Fix agent i and the reports t_{-i} of the other agents. Consider the corresponding allocation graph G_f. For any pair of types $s, t \in T$ such that $f(t) = f(s) = a$ for some $a \in A$, the payments must be equal, i.e. $\pi(t) = \pi(s) = \pi_a$. Therefore, π assigns a real number to every node in the graph. Incentive compatibility implies for any two outcomes $a, b \in A$ and all $t \in f^{-1}(b)$ that $v(b, t) - \pi_b \geq v(a, t) - \pi_a$, hence, $\pi_b \leq \pi_a + \ell_{ab}$.

For the other direction, define the payment π for agent i as follows. For any report vector of the other agents t_{-i}, consider the corresponding allocation graph G_f and fix a node potential p.

[6] We assume that arc lengths are strictly larger than $-\infty$. For allocation rules implementable in dominant strategies this is no restriction, as the incentive compatibility constraints imply finiteness of the arc lengths.

[7] Clearly, type and allocation graph depend on the agent i and reports t_{-i} of the other agents. In order to keep notation simple, we suppress the dependence on i and t_{-i} and will simply write T_f and G_f.

At aggregate report vector (t_i, t_{-i}) with outcome $a = f(t_i, t_{-i})$, let the payment be $\pi_a := p(a)$. Incentive compatibility now follows from the fact that p is a node potential in G_f, similarly to the above. □

Let

$$dist_{G_f}(a, b) = \inf_{P \in \mathcal{P}(a,b)} length(P).$$

In general, $dist_{G_f}(a, b)$ could be unbounded. However, if G_f does not contain a negative cycle (the *nonnegative cycle property*), then $dist_{G_f}(a, b)$ is finite, since the length of any (a, b)-path is lower bounded by $-\ell_{ba}$.

Lemma 1 *Fix an agent and some report vector of the other agents. The corresponding allocation graph G_f has a node potential if and only if it satisfies the nonnegative cycle property.*

Proof. Proofs can be found e.g. in Schrijver (2003) for finite A and in Rochet (1987) for infinite A. For completeness, we give a simple proof. If G_f has no negative cycle, then for any $a \in A$, $dist_{G_f}(a, \cdot)$ is well-defined, i.e. takes only finite values. The distances $dist_{G_f}(a, \cdot)$ define a node potential, because $dist_{G_f}(a, x) \leq dist_{G_f}(a, y) + \ell_{yx}$ for all $x, y \in A$. On the other hand, given a node potential p, add up the inequalities $p(y) - p(x) \leq \ell_{xy}$ for all arcs (x, y) on a cycle to prove that the cycle has nonnegative length. □

Observation 1 together with Lemma 1 yields a characterization of allocation rules that are implementable in dominant strategies (see also e.g. Rochet (1987)).

Observation 2 *The allocation rule f is implementable in dominant strategies if and only if all allocation graphs G_f obtained from a combination of an agent and a report vector of the other agents satisfy the nonnegative cycle property.*

From Lemma 1 and observations 1 and 2 it follows that for any allocation rule f implementable in dominant strategies, there exist node potentials in all allocation graphs G_f. The allocation rule f satisfies revenue equivalence if and only if in each allocation graph G_f, the node potential is uniquely defined up to a constant. Our main result states that this is the case if and only if distances are anti-symmetric in every G_f.

Theorem 1 (Characterization of Revenue Equivalence) *Let f be an allocation rule implementable in dominant strategies. f satisfies revenue equivalence if and only if in all allocation graphs G_f obtained from a combination of an agent and a report vector of the other agents, distances are anti-symmetric, i.e. $dist_{G_f}(a, b) = -dist_{G_f}(b, a)$ for all $a, b \in A$.*

Proof. Suppose first that f satisfies revenue equivalence. Fix a combination of an agent and a report vector of the other agents and let G_f be the corresponding allocation graph. Let $a, b \in A$. The functions $dist_{G_f}(a, \cdot)$ and $dist_{G_f}(b, \cdot)$ are node potentials in G_f. As any two node potentials differ only by a constant, we have that $dist_{G_f}(a,\cdot)-dist_{G_f}(b,\cdot)$ is a constant function. Especially, for a and b we get that $dist_{G_f}(a,a)-dist_{G_f}(b,a)=dist_{G_f}(a,b)-dist_{G_f}(b,b)$. Clearly, $dist_{G_f}(a,a)=dist_{G_f}(b,b)=0$ and hence $dist_{G_f}(a,b)=-dist_{G_f}(b,a)$.

Now suppose that $dist_{G_f}(a,b)=-dist_{G_f}(b,a)$ for all $a,b\in A$. Let $a,b\in A$. Let P_{ab} be an (a,b)-path with nodes $a=a_0,a_1,\ldots,a_k=b$. For any node potential p add up the inequalities $p(a_i)-p(a_{i-1})\leq \ell_{a_{i-1}a_i}$ for $i=1,\ldots,k$. This yields $p(b)-p(a)\leq length(P_{ab})$. Therefore,

$$p(b)-p(a)\leq \inf_{P\in\mathcal{P}(a,b)} length(P)=dist_{G_f}(a,b).$$

Similarly, $p(a)-p(b)\leq dist_{G_f}(b,a)$. Therefore, $-dist_{G_f}(b,a)\leq p(b)-p(a)\leq dist_{G_f}(a,b)$. Since $dist_{G_f}(a,b)=-dist_{G_f}(b,a)$, $p(b)-p(a)=dist_{G_f}(a,b)$ for any node potential p. Hence, any potential is completely defined, once $p(a)$ has been chosen for some outcome a. Thus, any two node potentials can only differ by a constant and f satisfies revenue equivalence. □

An analogous characterization holds for type graphs as well. One can check that all previous arguments still apply when using type graphs. On the other hand, note the following relation of node potentials in G_f and node potentials in T_f. Given a node potential p^G in G_f, we can define a node potential p^T in T_f by letting $p^T(t):=p^G(f(t))$ for any $t\in T$. In fact, let ℓ^G and ℓ^T denote the arc lengths in G_f and T_f respectively and observe that $\ell^G_{ab}=\inf\{\ell^T_{st}\mid s\in f^{-1}(a),\ t\in f^{-1}(b)\}$. Then, for any $s,t\in T$, $p^T(t)=p^G(f(t))\leq p^G(f(s))+\ell^G_{f(s)f(t)}\leq p^T(s)+\ell^T_{st}$ and p^T is a node potential. On the other hand, given a node potential p^T in T_f, let $p^G(a):=p^T(s)$ for any $s\in f^{-1}(a)$. Note that p^G is well-defined as $f(s)=f(t)=a$ implies $\ell^T_{st}=0$ and hence $p^T(s)=p^T(t)$. Furthermore, for any $a,b\in A$ and any $s\in f^{-1}(a),t\in f^{-1}(b)$, $p^G(a)=p^T(s)\leq p^T(t)+\ell^T_{ts}=p^G(b)+\ell^T_{ts}$. Hence, $p^G(a)\leq p^G(b)+\ell^G_{ba}$ and p^G is a node potential in G_f. Consequently, there is a one-to-one relationship between node potentials in G_f and node potentials in T_f. This insight together with a proof similar to the one of Theorem 1 yield the following corollary.

Corollary 1 (Characterization of Revenue Equivalence on Type Graphs) *Let f be an allocation rule that is implementable in dominant strategies. Then f satisfies revenue equivalence if and only if in all type graphs T_f obtained from a combination of an agent and a report vector of the other agents, distances are anti-symmetric, i.e. $dist_{T_f}(s,t)=-dist_{T_f}(t,s)$ for all $s,t\in T$.*

4 Discussion

In settings with multi-dimensional type spaces, finite A, and valuation functions that are linear in types, implementability implies that the allocation rule f, viewed from a single agent perspective as a vector field that maps multi-dimensional types on lotteries over outcomes, has a potential function F. One can easily verify that this property has the following interpretation on type graphs: the length of a shortest path in T_f from some type s to some type t is upper bounded by a path integral of the vector field f, or equivalently $F(t) - F(s)$. From this it follows easily that $dist_{T_f}(s,t) = -dist_{T_f}(t,s)$, i.e. revenue equivalence holds. In particular, $dist_{T_f}(s,t) = F(t) - F(s)$ for any potential function F. This connection between implementability, potential functions and revenue equivalence is also established in Jehiel, Moldovanu, and Stacchetti (1996), Jehiel, Moldovanu, and Stacchetti (1999), Jehiel and Moldovanu (2001) and Krishna and Maenner (2001).

It is interesting to compare our result with the characterization by Chung and Olszewski (2007). First, we introduce the notation used by Chung and Olszewski (2007) and restate their characterization theorem. Let A be countable. As before, regard everything from the perspective of a single agent. Let A_1, A_2 be disjoint subsets of A and $r : A_1 \cup A_2 \to \mathbb{R}$. For every $\varepsilon > 0$, let

$$\mathcal{T}_1(\varepsilon) = \bigcup_{a_1 \in A_1} \{t \in T \mid \forall a_2 \in A_2 : v(a_1,t) - v(a_2,t) > r(a_1) - r(a_2) + \varepsilon\}$$

and

$$\mathcal{T}_2(\varepsilon) = \bigcup_{a_2 \in A_2} \{t \in T \mid \forall a_1 \in A_1 : v(a_1,t) - v(a_2,t) < r(a_1) - r(a_2) - \varepsilon\}.$$

Finally, let $\mathcal{T}_i = \cup_{\varepsilon>0}\mathcal{T}_i(\varepsilon)$, $i = 1,2$. Observe that $\mathcal{T}_1 \cap \mathcal{T}_2 = \emptyset$. Call the type space T *splittable* if there are A_1, A_2 and r such that $T = \mathcal{T}_1 \cup \mathcal{T}_2$ and $\mathcal{T}_i \neq \emptyset$ for $i = 1,2$.

Theorem 2 (Chung and Olszewski (2007)) *If A is finite, the following two statements are equivalent.*

(i) All f that are implementable in dominant strategies satisfy revenue equivalence.

(ii) For all agents T_i is not splittable.

If A is not finite, but countable, (ii) implies (i).

To see the connection between the allocation graph defined in Section 3 and the notion of splittable, we outline a proof that $(ii) \Rightarrow (i)$. Suppose an allocation rule f implementable in dominant strategies that fails revenue equivalence. Since f is implementable, the allocation graphs satisfy the non-negative cycle property. Since revenue equivalence is violated, Theorem 1 implies an agent i and reports of the other agents t_{-i} such that in the corresponding allocation graph G_f, $dist_{G_f}(a^*,b^*) + dist_{G_f}(b^*,a^*) > 0$ for some $a^*, b^* \in A$. Assume the perspective of agent i. We show that this implies that T_i is splittable.

Define $d(a) = dist_{G_f}(a^*,a) + dist_{G_f}(a,a^*)$ for all $a \in A$. Since the function d takes only countably many values, there exists $z \in \mathbb{R}$ such that the following sets form a non-trivial partition of A: $A_1 = \{a \in A \mid d(a) > z\}$, $A_2 = \{a \in A \mid d(a) < z\}$. Observe that for every $a_1 \in A_1$, there exists $\varepsilon(a_1) > 0$ such that $d(a_1) > z + \varepsilon(a_1)$. Similarly, for every $a_2 \in A_2$, there exists $\varepsilon(a_2) > 0$

such that $d(a_2) < z - \varepsilon(a_2)$. It is now straightforward to verify that the sets A_i 'split' the type space.

Notice that in Theorem 1 no assumption on the cardinality of A is made, whereas in Theorem 2, A is assumed finite or countable. On the other hand, Theorem 1 imposes a condition on the allocation rule, whereas Theorem 2 characterizes T and v such that *all* allocation rules that are implementable in dominant strategies satisfy revenue equivalence. The principle difference between these settings is illustrated by the following example.

A principal has one unit of a perfectly divisible good to be distributed among n agents. The type of agent i is his demand $t_i \in (0, 1]$. Given the reports $t \in (0, 1]^n$ of all agents, an allocation rule $f\colon (0, 1]^n \to [0, 1]^n$ assigns a fraction of the good to every agent such that $\sum_{i=1}^n f_i(t) \leq 1$. If an agent's demand is met, he incurs a disutility of 0, otherwise his disutility is linear in the amount of unmet demand. More precisely, agent i's valuation if he is assigned quantity q_i is

$$v_i(q_i, t_i) = \begin{cases} 0, & \text{if } q_i \geq t_i, \\ q_i - t_i, & \text{if } q_i < t_i. \end{cases}$$

In this context, payments are reimbursements from the principal for unmet demand. This valuation function appears in Holmström (1979) as an example to demonstrate that his smooth path-connectedness assumption cannot be weakened. Likewise, the example can be used to show that the convexity assumption of the valuation function in Krishna and Maenner (2001) cannot be relaxed.

Call an allocation rule f *dictatorial*, if there is an agent i that always gets precisely his demanded quantity, $f_i(t_i, t_{-i}) = t_i$ for all t_{-i}. This rule is implementable, and as shown in Holmström (1979), fails revenue equivalence. However there are implementable rules in this setting that satisfy revenue equivalence:

Theorem 3 *For the demand rationing problem, the proportional allocation rule f with $f_i(t) = t_i / \sum_{j=1}^n t_j$ for $i = 1, \ldots, n$ is implementable and satisfies revenue equivalence.*

The proof uses the type graph for any agent i and any fixed report t_{-i} of other agents and verifies implementability by using Lemma 1. A fixed point argument is used to show that the distance function is anti-symmetric. Thus revenue equivalence holds due to Theorem 1.[8]

As this setting of T and v allows for allocation rules that satisfy revenue equivalence as well as for rules that don't, any theorem describing sufficient conditions for *all* implementable f to satisfy revenue equivalence is necessarily silent.

[8] A complete proof can be found in Heydenreich, Müller, Uetz, and Vohra (2008).

Acknowledgements

We thank the editor and 3 anonymous referees for very useful comments and suggestions. We also thank participants of Dagstuhl Seminars 07271 and 07471 for their valuable feedback.

References

Chung, K.-S. and W. Olszewski (2007). A non-differentiable approach to revenue equivalence. *Theoretical Economics 2*, 469–487.

Green, J. and J.-J. Laffont (1977). Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica 45*(2), 427–438.

Gui, H., R. Müller, and R. Vohra (2004). Dominant strategy mechanisms with multidimensional types. Discussion Paper 1392, The Center for Mathematical Studies in Economics & Management Sciences, Northwestern University, Evanston, IL.

Heydenreich, B., R. Müller, M. Uetz, and R. Vohra (2008). Characterization of revenue equivalence. Meteor Research Memorandum RM/08/01, Maastricht University.

Holmström, B. (1979). Groves' scheme on restricted domains. *Econometrica 47*(5), 1137–1144.

Jehiel, P. and B. Moldovanu (2001). Efficient design with interdependent valuations. *Econometrica 69*(5), 12371259.

Jehiel, P., B. Moldovanu, and E. Stacchetti (1996). How (not) to sell nuclear weapons. *American Economic Review 86*(4), 814–829.

Jehiel, P., B. Moldovanu, and E. Stacchetti (1999). Multidimensional mechanism design for auctions with externalities. *Journal of Economic Theory 85*(2), 258–293.

Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys 13*(3), 227–286.

Krishna, V. (2002). *Auction Theory*. Academic Press.

Krishna, V. and E. Maenner (2001). Convex potentials with an application to mechanism design. *Econometrica 69*(4), 1113–1119.

Milgrom, P. (2004). *Putting Auction Theory to Work*. New York: Cambridge University Press.

Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica 70*(2), 583–601.

Müller, R., A. Perea, and S. Wolf (2007). Weak monotonicity and Bayes-Nash incentive compatibility. *Games and Economic Behavior 61*(2), 344–358.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research 6*(1), 58–73.

Rochet, J.-C. (1987). A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics 16*(2), 191–200.

Saks, M. and L. Yu (2005). Weak monotonicity suffices for truthfulness on convex domains. In *Proc. 6th ACM conference on Electronic Commerce*, pp. 286 – 293.

Schrijver, A. (2003). *Combinatorial Optimization*, Volume A. Springer.

Suijs, J. (1996). On incentive compatibility and budget balancedness in public decision making. *Economic Design 2*, 193–209.