Challenges in Algorithmic Mechanism Design

THÈSE N^{O} 5711 (2013)

PRÉSENTÉE LE 26 AVRIL 2013

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS
LABORATOIRE DE THÉORIE ET APPLICATION D'ALGORITHMES
PROGRAMME DOCTORAL EN INFORMATIQUE, COMMUNICATIONS ET INFORMATION

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Paul David DÜTTING

acceptée sur proposition du jury:

Prof. F. Eisenbrand, président du jury
Prof. M. Henzinger, Prof. B. Faltings, directeurs de thèse
Prof. T. Hofmann, rapporteur
Prof. D. C. Parkes, rapporteur
Prof. B. Vöcking, rapporteur

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2013

To Ina

Acknowledgements

To my thesis director Monika Henzinger and my thesis co-director Boi Faltings, for being great supervisors. To Friedrich Eisenbrand, for agreeing to act as thesis jury president. To Thomas Hofmann, David C. Parkes, and Berthold Vöcking, for agreeing to be part of my thesis jury. To Rüdiger Urbanke, for supporting me in administrative matters. To Felix Fischer, Monika Henzinger, Pichayut Jirapinyo, Radu Jurca, John K. Lai, Benjamin Lubin, David. C. Parkes, Martin Starnberger, and Ingmar Weber, for exciting collaborations. To Radu Jurca and David C. Parkes, for wonderful internships at Google and Harvard. To Jean-Cédric Chappelier, Felix Fischer, Florent Garcin, Monika Henzinger, Veronika Loitzenbauer, David C. Parkes, and Martin Starnberger, for valuable feedback on parts of this thesis. And, last but not least, to my wife, my family, and my friends, for their help and support.

Abstract

This thesis addresses three challenges in algorithmic mechanism design, which seeks to devise computationally efficient mechanisms consisting of an outcome rule and a payment rule that implement desirable outcomes in strategic equilibrium.

The first challenge that we address is the *design of expressive mechanisms*, i.e., mechanisms that allow the participating agents to express rich preferences. We focus on multi-item auctions with unit demand. For this setting we present the most expressive polynomial-time mechanism known to date that is incentive compatible for non-degenerate inputs. This mechanism can, e.g., be used in ad auctions with per-click and per-impression valuations and it can handle a large variety of soft and hard budget constraints.

The second challenge that we consider is the *analysis of simplicity-expressiveness tradeoffs*. We develop tools for analyzing how simplification, i.e., restricting the message space, affects the set of equilibria of a mechanism. We use these tools to analyze two representative settings, sponsored search auctions and combinatorial auctions. We find that in both cases simplification can be beneficial, either by ruling out undesirable equilibria or by promoting desirable ones. In the case of sponsored search auctions our analysis leads to a strong argument in favor of the mechanism that is used by all major search engines.

Finally, and this is the third challenge, we consider the *design of approximately strategyproof mechanisms*. We present a framework that exploits a remarkably close connection between discriminant-based classification and the design of strategyproof mechanisms. For a given algorithmically specified outcome rule our framework finds a payment rule that makes the resulting mechanism maximally strategyproof. We support our theoretical findings by applying our framework to a multi-minded combinatorial auction with a greedy outcome rule and to an assignment problem with egalitarian outcome rule.

Keywords

Algorithmic game theory, algorithmic mechanism design, expressiveness, simplicity, approximate strategyproofness

Résumé

Cette thèse aborde trois défis dans la théorie algorithmique des mécanismes d'incitation dont le but est de concevoir des mécanismes computationnels efficaces qui implémentent des résultats souhaités dans un équilibre stratégique.

Le premier défi que nous abordons est la *conception de mécanismes expressifs*, c'est-à-dire des mécanismes qui permettent d'exprimer des préférences riches. Nous nous concentrons sur les enchères à plusieurs articles avec demande unitaire. Dans ce cadre, nous présentons le mécanisme le plus expressif en temps polynomial connu à ce jour qui soit compatible avec les mesures d'incitation pour des entrées non-dégénérées. Ce mécanisme peut être utilisé dans des enchères pour des emplacements publicitaires avec valorisation par click ou par impression et il peut gérer une large variété de contraintes budgétaires.

Le deuxième défi que nous considérons est *l'analyse du compromis simplicité-expressivité*. Nous développons des outils pour analyser comment la simplification, c'est-à-dire la restriction de l'espace des messages, influence l'ensemble des équilibres d'un mécanisme. Nous utilisons ces outils pour analyser les enchères de recherche sponsorisée et les enchères combinatoires. Nous montrons que dans les deux cas la simplification peut être bénéfique, soit en excluant les équilibres non-désirables soit en promouvant ceux qui sont souhaités. Dans le cas des enchères de recherche sponsorisée, notre analyse mène à un argument fort en faveur du mécanisme utilisé par la plupart des moteurs de recherche.

Finalement, le troisième défi que nous considérons concerne la *conception de mécanismes d'incitation approximative*. Nous présentons une méthode qui exploite une similarité remarquable entre la classification discriminante et la conception des mécanismes d'incitation. Pour une règle de résultat donnée de façon algorithmique, notre méthode trouve une règle de paiement qui rend le mécanisme résultant maximalement compatible avec les mesures d'incitation. Nous renforçons notre conclusion théorique en appliquant la méthode à une enchère combinatoire à intentions multiples avec une règle de résultat « gloutonne » ainsi qu'à un problème d'attribution avec une règle de résultat égalitaire.

Mot-Clés

Théorie algorithmique des jeux, théorie algorithmique des mécanismes d'incitation, simplicité, expressivité, incitation approximative

Zusammenfassung

Diese Dissertation befasst sich mit drei Herausforderungen des algorithmischen Mechanismenentwurfs, dessen Ziel effizient berechenbare Mechanismen mit erstrebenswerten strategischen Gleichgewichten sind.

Die erste Herausforderung ist der *Entwurf ausdrucksstarker Mechanismen*, welche die Spezifikation vielfältiger Präferenzen ermöglichen. Der Fokus liegt dabei auf Versteigerungen mehrerer Gegenstände, in denen jeder Bieter höchstens einen Gegenstand erhalten kann. Es wird ein effizient berechenbarer Mechanismus für bisher nicht zu bewältigende Präferenzen entwickelt, der ausser für degenerierte Eingaben nicht manipulierbar ist. Dieser Mechanismus kann zur Versteigerung von Werbeflächen im Internet verwendet werden, bei der die Bieter teilweise an Klicks und teilweise an der Häufigkeit der Darstellung interessiert sind. Ausserdem ermöglicht er die Spezifikation diverser Budgetbeschränkungen.

Die zweite Herausforderung ist das *Abwägen von Vor- und Nachteilen der Ausdruckstärke*. Es werden Techniken entwickelt, um die Auswirkung von Vereinfachung (Einschränkung der erlaubten Gebote) auf die Gleichgewichte eines Mechanismus zu analysieren. Die Techniken werden auf zwei Probleme angewendet, Versteigerungen von Suchergebnissen und kombinatorische Auktionen. In beiden Fällen zeigt die Analyse, dass Vereinfachung hilfreich sein kann, entweder durch den Ausschluss schlechter oder die Begünstigung guter Gleichgewichte. Im ersten Fall führt die Analyse zu einem überzeugenden Argument für den Mechanismus, der von allen grösseren Suchmaschinenbetreibern verwendet wird.

Schliesslich und das ist die dritte Herausforderung, wird der *Entwurf von annähernd wahrheitsgemässen Mechanismen* betrachtet. Es wird eine Methode vorgestellt, die einen überraschenden Zusammenhang zwischen diskriminantenbasierter Klassifikation einerseits und dem Entwurf wahrheitsgemässer Mechanismen andererseits ausnutzt. Diese Methode findet zu einer algorithmisch spezifizierten Ausgaberegel automatisch eine Zahlungsregel, die den resultierenden Mechanismus so wahrheitsgemäss wie möglich macht. Die Methode wird auf kombinatorische Auktionen mit einer Heuristik zur Maximierung des Gemeinwohls und auf ein Zuweisungsproblem mit egalitärer Ausgaberegel angewendet.

Schlüsselwörter

Algorithmische Spieltheorie, algorithmischer Mechanismenentwurf, Ausdrucksstärke, Vereinfachung, annähernd wahrheitsgemässe Mechanismen

Contents

List of Figures

List of Tables

Introduction

Computer science and *game theory* developed almost at the same time, namely in the 1950s. And some of the most important researchers at that time, such as von Neumann, contributed to computer science *and* game theory. Still, somewhat surprisingly, both disciplines developed more or less separately from each other after their conception. This did not change until the emergence and popularization of the *Internet* in the 1990s. As of today computer science and game theory are intertwined again, and problems at the intersection of the two fields are an active area of research.

An important area of computer science is *algorithm design*, which aims at devising step-by-step descriptions, or algorithms, for how given a certain input a certain output can be computed. A representative problem is sorting, where a given set of numbers has to be sorted in, say, increasing order. A main concern of algorithm design is *computational efficiency*. For this it analyzes how many steps an algorithm takes for an input of a given size, in the worst case over all inputs with that size, until it has found the output. Of particular interest are so-called polynomial-time algorithms, for which the number of steps is bounded by a polynomial in the input size.

An important branch of game theory is *mechanism design*. The main object of study of mechanism design are mechanisms that consist of an outcome rule and a payment rule. The input to both is provided by agents, who employ preferences over outcomes and payments, and can misreport the input if they find it beneficial to do so. A canonical example is the auction of a single item, where the input are the values that the agents have for the item, the agents' utility is value minus price paid if they win the item and zero otherwise, and the objective is to assign the item to the agent with the highest value. The goal of mechanism design is more generally to implement desirable outcomes taking the *strategic behavior* of the agents into account, which typically involves incentivizing the agents to truthfully report their preferences.

In this thesis we address three challenges in *algorithmic mechanism design*, which—just as algorithm design—insists on polynomial-time computability and—just as mechanism design—works under the premise that the input is provided by selfish agents and therefore seeks to incentivize agents to report truthfully.

Challenge 1: Design of Expressive Mechanisms

The first challenge that we address is the *design of expressive mechanisms*, i.e., mechanisms that allow agents to express rich preferences via their utility functions. Motivated by the ever increasing demand for expressive auctions in applications on the Web, such as Google's and Microsoft's ad auctions or auctions on platforms such as eBay, we focus on the domain of *multi-item auctions with unit demand*.

The problem solved by these auctions is essentially a matching and a pricing problem in which the goal is to find a bidder optimal, envy free solution. A solution to this problem is *envy free* if every bidder is assigned an item that maximizes his utility at the current prices, and it is *bidder optimal* if it gives each bidder the highest utility among all envy free solutions. From an economic point of view a bidder optimal, envy free solution is desirable as it corresponds to the *competitive equilibrium* with *maximum payoffs* to the bidders. Maximizing the payoffs to the bidders is also a reasonable design goal for the auctioneer as it guarantees participation and thus revenue in the long run.

Standard mechanisms for this problem typically restrict the bidders' ability to express their preferences by forcing them to submit utility functions that are (1) linear, with identical slopes across the items, and (2) continuous in the price. More recently, a mechanism was proposed that adds a single discontinuity per bidder-item pair. Still, as we show, these mechanisms cannot be applied to many problems, including ad auctions in which some of the bidders have per-click valuations and others have per-impression valuations or auctions on eBay in which the bidders have soft *and* hard budgets.

We overcome these limitations by presenting the first *polynomial-time mechanism* for computing a bidder optimal, envy free solution for *piece-wise linear* utility functions with *non-identical slopes* and *multiple discontinuities*. We also analyze under which conditions mechanisms that compute a bidder optimal, envy free solution incentivize truthful reporting. We show that under a certain non-degeneracy assumption regarding the input no bidder can gain by misreporting his utility functions, no matter what the other bidders report. Mechanisms with this property are *incentive compatible* (or *strategyproof*).

To summarize, we present the *most expressive polynomial-time mechanism* for the problem of finding a bidder optimal, envy free solution in multi-item auctions with unit demand *that is incentive compatible for non-degenerate inputs*.

Challenge 2: Analysis of Simplicity-Expressiveness Tradeoffs

The second challenge that we address is the *analysis of simplicity-expressiveness tradeoffs*. We contribute to this challenge by analyzing *simplified mechanisms*, i.e., mechanisms that are derived from another mechanism by restricting the ability of the agents to express their preferences. We refer to these restrictions as restrictions of the message space. The study of simplified mechanisms is motivated by the fact that in many practical situations truthful direct-revelation mechanisms, in which the agents truthfully reveal their preferences, are either unattainable or undesirable. They may be unattainable because of computational limitations, and they may be undesirable because of the existence of other, non-truthful, equilibria with undesirable economic properties.

We develop tools that enable the analysis of how simplifications, i.e., restrictions of the message space, affect the set of equilibria of a mechanism. One important property in this context is *tightness*, which requires that the restriction does not introduce new equilibria. Tightness is desirable as it precludes the introduction of new and potentially bad equilibria. Orthogonal to tightness is *totality*, which requires that the restriction preserves all equilibria. To the end of equilibrium selection, however, totality needs to be relaxed. A meaningful relaxation that we consider is the preservation of desirable equilibria.

We apply these tools to *two representative settings*, sponsored search auctions and combinatorial auctions, each being a canonical example for complete information and incomplete information analysis, respectively. For *sponsored search auctions* we observe that expressive versions of the standard mechanisms for this problem always permit an efficient, zero revenue equilibrium, while the restrictions on the message space used in practice are tight, preserve desirable equilibria, and rule out zero revenue equilibria. This shows that these restrictions strictly improve the set of equilibria. We also show that the mechanism used in practice guarantees the existence of a desirable equilibrium and good revenue in all equilibria for a broader class of inputs than alternative mechanisms. For *combinatorial auctions* we characterize precisely which simplifications of the standard mechanism for this problem are tight, and as such guarantee that the worst equilibrium of the simplified mechanism is no worse than the worst equilibrium of the original mechanism. We also show that each of these simplifications provides a different focal (truthful) equilibrium, and that this can have a significant impact on social welfare at equilibrium.

We also observe that the *amount of information available to the agents* plays an important role for the tradeoff between simplicity and expressiveness. In the sponsored search setting, both the existence of a zero revenue equilibrium in the expressive mechanism, and the existence of a desirable equilibrium in the simplified mechanism, rely on the assumption of complete information. For combinatorial auctions, there exist other tight simplifications provided that the agents have complete information.

To summarize we develop a *toolbox for the analysis of simplified mechanisms* and apply it to sponsored search auctions and combinatorial auctions. We show that in both cases *simplification can be beneficial*, either by precluding bad or by promoting good equilibria. In the case of sponsored search auctions our analysis leads to a *strong argument in favor of the mechanism that is used by all major search engines.*

Challenge 3: Design of Approximately Strategyproof Mechanisms

The third and last challenge that we address is the *design of approximately strategyproof mechanisms.* Adopting approximate strategyproofness as a design goal is motivated by the fact that the classical approach of deriving the optimal mechanism subject to exact strategyproofness is associated with several disadvantages. First, it requires a *de novo* design for each domain and in some domains it can be analytically cumbersome to derive the optimal mechanism. Second, adopting incentive compatibility as a hard constraint may preclude mechanisms with useful economic properties. Third, the optimal mechanism may have an outcome rule or payment rule that is computationally intractable.

The notion of approximate strategyproofness that we adopt is *minimization of expected ex post regret.* The ex post regret an agent has for truthful reporting in a given instance is the maximum amount by which his utility could be increased through a misreport holding the reports of others fixed. The expected ex post regret is the average ex post regret over all agents and all preferences, calculated with respect to a distribution on preferences. The expected ex post regret quantifies the potential gain from manipulation. A mechanism with zero expected ex post regret is strategyproof.

By replacing the incentive compatibility requirement with the goal of minimizing expected ex post regret, we are able to adapt *statistical machine learning techniques* to the design of payment rules. Specifically, given an algorithmically specified outcome rule and a distribution over preferences we train a discriminant-based classifier with a special structure to predict the outcome rule. We then use the learned discriminant to define a payment rule. We show that if the discriminant-based classifier is *exact*, then the resulting mechanism is *strategyproof.* We also show that if the discriminant-based classifier *minimizes classification error*, then the resulting mechanism *minimizes expected ex post regret.*

We support our theoretical findings by applying our framework to *two applications*, a multi-minded combinatorial auction with a greedy outcome rule and an assignment problem with an outcome rule that maximizes egalitarian welfare. For the former it can be shown that no payment rule exists that makes it strategyproof, and for the latter it is not clear whether such a payment rule exists. Our experiments are encouraging in that they show that in both cases our framework produces mechanisms with low expected ex post regret.

To conclude, we propose a *new paradigm for the design of approximately strategyproof mechanisms* that exploits a remarkably direct connection between discriminant-based classification and strategyproof mechanism design. Specifically, given an algorithmically specified outcome rule our framework *automatically finds a payment rule* that makes the resulting mechanism *maximally strategyproof.*

Organization

The rest of this thesis is organized as follows. We present our results regarding the first challenge in Chapter 1, our results regarding the second challenge in Chapter 2, and our results regarding the third challenge in Chapter 3. We conclude with some remarks on the challenges treated in this thesis and directions for future work.

1 An Expressive Mechanism for Auctions on the Web

1.1 Introduction

Auctions are widely used on the Web. They are, e.g., used by Google and Microsoft for selling sponsored search results and on platforms such as eBay for selling a broad variety of goods and services. In these and in many other applications the auctions in use restrict the bidders to receive at most one item. In other words, the auctions in use are often *multi-item auctions with unit demand*.

The problem solved by these auctions is essentially a *matching* and *pricing* problem. In this problem n bidders have to be matched to k items. Each bidder i has a utility function $u_{i,j}(p_j)$ that expresses his utility for being matched to item j at price p_j. A solution (μ, p) consisting of a matching μ and prices p is said to be *envy free* if at the current prices every bidder (weakly) prefers the item that he is currently matched to over every other item. An envy free solution (μ, p) is said to be *bidder optimal* if the utility of every bidder is at least as high as in every other envy free solution (μ', p').

From an economic point of view an envy free solution in which the price of every unsold item is identical to the item's reserve price is desirable because it represents a *competitive equilibrium*.[1] Under the additional requirement on the prices of unsold items the bidder optimal, envy free solution is the competitive equilibrium with *maximum payoffs* to the bidders. Maximizing the payoffs to the bidders is also a reasonable objective for the auctioneer as it guarantees participation and thus revenue in the long run.

1.1.1 Limitations of Current Mechanisms

Standard mechanisms for auctions on the Web such as the *Vickrey-Clarke-Groves (VCG)* mechanism [27, 9, 17] or the *Generalized Second Price (GSP)* mechanism [15, 25] nicely fit into

[1] The additional requirement regarding the prices of unsold items is needed to ensure that the auctioneer prefers not to sell these items.

the above model. For linear utilities of the form $u_{i,j}(p_j) = v_{i,j} - p_j$, where $v_{i,j}$ is a bidder-item dependent valuation, the VCG mechanism finds a bidder optimal, envy free solution [21]. (Locally) envy free equilibria of the GSP mechanism for linear utilities of the form $u_{i,j}(p_j) = \alpha_j \cdot v_i - p_j$, where α_j is an item-dependent constant and v_i is an agent-dependent valuation, are discussed in [15, 25]. A main drawback of these mechanisms, however, is the *limited expressiveness* that they offer to the bidders: (1) They typically restrict the utilities to be *linear* in the price, with *identical slopes* across the items. (2) They usually require the utilities to be *continuous* in the price.

The *General Auction Mechanism (GAM)* of Aggarwal et al. [1] takes a first step towards addressing the problem of limited expressiveness. It applies to *linear* utilities with *identical slopes* and a *single discontinuity* per bidder-item pair. More specifically, the utilities are of the form $u_{i,j}(p_j) = v_{i,j} - p_j$ if $p_j \leq m_{i,j}$ and $u_{i,j}(p_j) = -\infty$ otherwise, where $m_{i,j}$ is a so-called bidder-item dependent *maximum price*. GAM requires the input to be in *general position*. For inputs in general position it finds a bidder optimal, envy free solution in polynomial time. It also preserves another desirable property of the original model. Namely, no bidder can misreport his valuations and/or maximum prices to achieve a higher utility. Mechanisms with this property are *incentive compatible* (or *strategyproof*).

The general position assumption mandates that in a certain weighted multi-graph defined on the basis of the input no two walks have exactly the same weight. As this is rather unlikely for generic weights, e.g., randomly generated ones, inputs that are not in general position can be regarded as degenerate.[2]

Despite its generality GAM has two major limitations: (1) It can only handle linear utility functions with *identical slopes.* (2) It can only handle a *single discontinuity* with a jump to $-\infty$ per bidder-item pair. We illustrate why and when these shortcomings are problematic by means of two examples.

Example 1. Per-click vs. per-impression valuations (This example motivates linear utilities with *non-identical* slopes.) Consider an ad auction with bidders with per-click valuations $v_{i,j}^{click}$ and bidders with per-impression valuations $v_{i,j}^{imp}$. The former are envy free if $u_i^{click} \geq v_{i,j}^{click} - p_j^{click}$ for all j and the latter are envy free if $u_i^{imp} \geq v_{i,j}^{imp} - p_j^{imp}$ for all j. Suppose that the mechanism collects per-click valuations and charges per-click prices. That is,

$$u_{i,j}(p_j^{click}) = v_{i,j}^{click} - p_j^{click}. \tag{1.1}$$

A bidder with per-impression valuations can translate his valuations into per-click valuations using the click trough rate $ctr_{i,j}$ as follows: $v_{i,j}^{click} = v_{i,j}^{imp}/ctr_{i,j}$. That is, he reports $u_{i,j}(p_j^{click}) = v_{i,j}^{imp}/ctr_{i,j} - p_j^{click}$. Now suppose that given the per-click valuations, the mech-

[2] General position can also be defined as algebraic independence meaning that no non-tautological equations relating the input values are satisfied. As we only require certain equations to be violated this alternate definition of general position is sufficient but not necessary for us.

anism computes an envy free solution (μ, p^{click}) consisting of a matching μ and per-click prices p^{click}. That is, for every matched bidder-item pair $(i, j) \in \mu$ and all items $j' \neq j$ we have:

$$v_{i,j}^{click} - p_j^{click} \geq v_{i,j'}^{click} - p_{j'}^{click}. \tag{1.2}$$

What we actually want for bidders i with per-impression valuations $v_{i,j}^{imp}$ is that for $(i, j) \in \mu$ and all $j' \neq j$:

$$v_{i,j}^{imp} - p_j^{imp} \geq v_{i,j'}^{imp} - p_{j'}^{imp} \tag{1.3}$$

But if we take (1.2), replace $v_{i,j}^{click}$ with $v_{i,j}^{imp}/ctr_{i,j}$, p_j^{click} with $p_j^{imp}/ctr_{i,j}$, and multiply by $ctr_{i,j}$, then we get

$$v_{i,j}^{imp} - p_j^{imp} \geq C \cdot (v_{i,j'}^{imp} - p_{j'}^{imp}), \tag{1.4}$$

where $C = ctr_{i,j}/ctr_{i,j'}$. That is, if $C < 1$, then (1.4) is *not* strong enough to guarantee envy freeness for per-impression bidders. With non-identical slopes this can be sidestepped by having bidders with per-impression valuations report

$$u_{i,j}(p_j^{click}) = v_{i,j}^{imp} - ctr_{i,j} \cdot p_j^{click}. \tag{1.5}$$

In this case the solution (μ, p^{click}) computed by the mechanism will be envy free for both types of bidders, i.e., the above problem does *not* arise.

Example 2. Soft and hard budgets (This example motivates *piece-wise linear* utilities with *non-identical slopes* and *multiple discontinuities.*) Suppose that bidder i wants to buy a car on eBay. In the current system it would be dangerous for bidder i to bid on more than one car at the same time due to the risk of winning and having to pay for several cars when one is already enough. With a GAM-like auction bidder i could bid on many cars at the same time while still being guaranteed that he gets at most one.

With expensive items, such as cars, bidder i's valuation $v_{i,j}$ for item j may exceed the amount c of cash that he possesses. In this case i might be willing to take out a loan if the price p_j of item j exceeds c. Assume that bidder i is offered a loan with a maximum amount of a, an interest rate of r, and a fixed fee of f. Then bidder i's utility for item j has the following form: (1) Because no interest is due for the first c dollars the utility function drops linearly with a slope of -1 from 0 to c. (2) At c it drops by the fixed fee f. (3) Afterwards, due to the interest rate r, every dollar spent causes $1 + r$ dollars in actual cost. Hence the utility function drops linearly with a slope of $-(1 + r)$ from c to $c + a$.

In addition to the soft budget constraint c, bidder i may have a hard budget constraint $b < c + a$, which is typically modeled by a jump to $-\infty$ at $p_j = b$. Hence bidder i's utility function for item j ultimately looks like this:

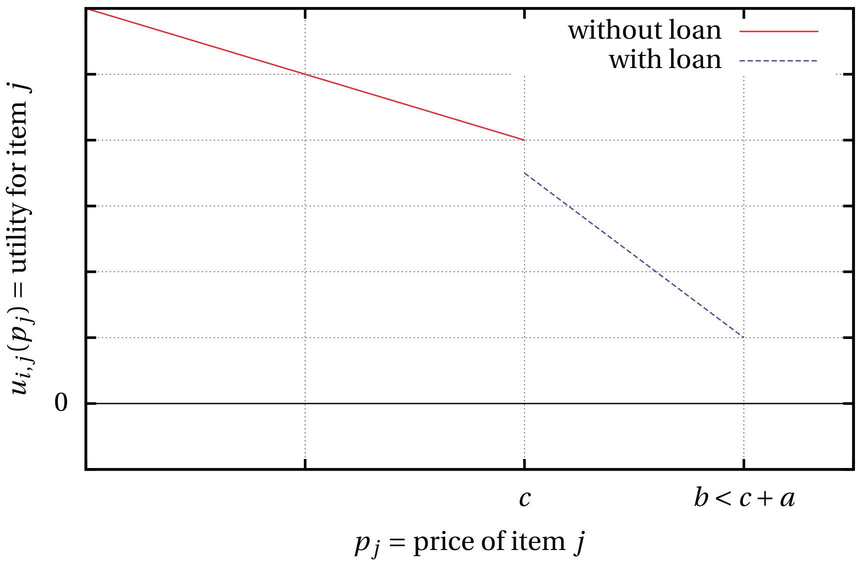

Figure 1.1: Piece-wise linear utility function with non-identical slopes and multiple discontinuities that arises from soft and hard budget constraints.

1.1.2 Our Contribution

We overcome the limitations of GAM by presenting a *polynomial-time mechanism* that computes a bidder optimal, envy free solution for *piece-wise linear* utilities with *non-identical slopes* and *multiple discontinuities*. Our mechanism is more expressive than GAM as it can be used in an ad auction to simultaneously auction off items to bidders with per-click and per-impression valuations (Example 1) and it can handle a large variety of soft and hard budget constraints (Example 2).

Our mechanism computes an envy free solution whose prices are minimal across all envy free solutions, which is sufficient for bidder optimality (Lemma 1). It starts by initializing the price of each item to the item's reserve price. Afterwards it tries to match one bidder after the other. For this it considers the *first choice graph*. The first choice graph consists of one node per bidder, one node per item, and an edge between a bidder and an item if the item is among the items that gives this bidder the highest utility. Given this graph our mechanism computes a *maximal alternating tree* whose root is the bidder that is to be matched. In a maximal alternating tree all paths from the root to a leaf alternate between unmatched and matched edges and none of these paths can be extended. If there is an alternating path from the bidder that is to be matched to an unmatched item, then the matching can be augmented by flipping the matched and unmatched edges along the alternating path (Lemma 3). Otherwise all items in the maximal alternating tree are *strictly overdemanded*, i.e., every subset of the items is wanted by strictly more bidders than there are items (Lemma 3). In this case our mechanism computes speeds at which the prices of these items need to be raised so that (i) all bidders that were matched to these items can still be matched to these items (possibly in a different way), (ii) there is a maximal alternating tree rooted at the unmatched bidder with respect to this matching that covers the same bidders and items, and (iii) the edges of the maximal alternating tree remain in the first choice graph. It can be shown that if the prices are raised according to these speeds then envy freeness and strict overdemand are preserved (Lemma 4 and Lemma 6). Our mechanism raises the prices according to these speeds until either (a) some bidder becomes interested in an item in which he was not interested before or (b) the end of a constant-slope interval of a utility function is reached. We show that if the prices were minimal before this update, then they are still minimal after this update (Lemma 7). To establish that this process finds a bidder optimal, envy free solution (Theorem 1) we argue as follows: The process terminates because (1) in updates according to (a) the maximal alternating tree under consideration grows which ensures that eventually there will be an alternating path along which the matching can be augmented and (2) updates according to (b) which can cause previously matched bidders to become unmatched can only occur a limited number of times. Envy freeness follows from termination because the matching is contained in the first choice graph at all times. For bidder optimality we argue inductively (using Lemma 7) that the prices are minimal.

We also analyze under which conditions on the input mechanisms that compute a bidder optimal, envy free solution for piece-wise linear utilities with non-identical slopes and multiple discontinuities are *incentive compatible*. We first provide an example that shows that *no* mechanism that computes a bidder optimal, envy free solution is incentive compatible *for all inputs*. We then show how to generalize the general position concept of Aggarwal et al. [1]. Finally, we prove that *every* mechanism that computes a bidder optimal, envy free solution is incentive compatible *for inputs in general position*.

Our main insight regarding incentive compatibility is that for inputs in general position at most one discontinuity is reached in each price increase of our mechanism (Lemma 9). In this case items that get unmatched due to a discontinuity in some iteration can get matched again in the subsequent iteration (Lemma 10). If our mechanism is biased towards matching previously matched items that got unmatched then it finds a bidder optimal, envy free solution (μ, p) in which (i) $p_j = r_j$ for all unmatched items j and (ii) $p_j = r_j$ for at least one matched item j (Proposition 1).[3] We use the existence of a bidder optimal, envy free solution with these properties to establish that (a) no solution can have higher utilities for *all* bidders (Lemma 11) and (b) if some solution gives higher utilities to *some* bidders, then there must be a bidder which does *not* get a higher utility that envies one of the bidders who gets a higher utility (Lemma 12). For incentive compatibility (Theorem 2) we then argue that, by (a), it is impossible that *all* bidders benefit from misreporting and, by (b), if only *some* bidders benefit from misreporting then at least one of the bidders who did *not* misreport is *not* envy free which contradicts the fact that the mechanism finds an envy free solution for the reported utility functions.

[3] Note that in this case the bidder optimal solution corresponds to the competitive equilibrium with maximum payoffs to the bidders.

1.1.3 Related Work

Continuous Utility Functions *Linear* utility functions with *identical* slopes were studied by Shapley and Shubik [23]. They formulated the problem of finding a matching that maximizes the social welfare as a linear program and observed that the dual program yields envy free prices. With the help of this formulation they proved the existence of a bidder optimal, envy free solution. Later Leonard [21] examined the incentives for misreporting the utility functions and found that the bidder optimal, envy free solution is identical to the solution of the VCG mechanism [27, 9, 17] and therefore incentive compatible. The "classic" mechanism for linear utilities with identical slopes is the *Multi-Item Auction* of Demange et al. [11], which is a variant of the *Hungarian Method* by Kuhn [20]. The basic idea of this mechanism is to start with prices all zero and to repeatedly raise the prices of overdemanded items by the same amount. This idea was generalized to *piece-wise linear* utility functions with *non-identical slopes* by Alkan [3, 4], who showed that the prices of overdemanded items need to be raised by different amounts.[4] The existence of a bidder optimal, envy free solution for general *non-linear* utilities was shown by Demange and Gale [10] using a lattice-theoretic argument. They also proved that any mechanism that finds a bidder optimal, envy free solution is incentive compatible. Recently, Alaei et al. [2] presented a novel, inductive characterization of the bidder optimal, envy free solution in this setting, which yields a constructive proof of existence. Although hardness results have been established for related problems (see, e.g., [12, 26]), it is not clear whether or under which conditions a bidder optimal, envy free solution can be found efficiently for such general utility functions.

Discontinuous Utility Functions *Linear* utilities with *identical slopes* and a *single discontinuity* were studied by Aggarwal et al. [1]. They gave a mechanism, which—for inputs in general position—is incentive compatible and finds a bidder optimal, envy free solution in polynomial time. Similar results to that of [1] were obtained by [5, 6] and [13]. In [19] it was shown how to find the smallest envy free prices for a *given* matching. Recently, Chen et al. [8] gave a polynomial-time mechanism for *consistent* utility functions. Note that all these results either assume *identical slopes* [1, 6, 13, 19], just a *single discontinuity* [1, 6, 13, 19, 8], or both. Also note that the utility functions that we study here are *not* consistent. The existence of a bidder optimal, envy free solution for more general, *non-linear* utility functions with *multiple discontinuities* was established in [14] using a lattice-theoretic argument. They also present conditions on the input under which every mechanism that computes a bidder optimal, envy free solution is incentive compatible. As in the continuous case, however, no polynomial-time mechanism for finding a bidder optimal, envy free solution for these general, non-linear utility functions is known.

[4] We adopt some of the techniques developed in [3, 4] for piece-wise linear utilities with non-identical slopes, but we adapt them to be able to cope with discontinuities and we also refine them to significantly improve upon the running time. In [3, 4] the running time is stated as $O(n^2 \cdot k^4 \cdot \prod_{i,j} t_{i,j})$, where $t_{i,j}$ is the number of constant-slope intervals of $u_{i,j}(\cdot)$.

1.2 Problem Statement

We are given a set I of n bidders and a set J of k items. The set of items J contains a dedicated dummy item that we denote j_0. For each bidder i we are given a constant o_i, called the *outside option*, which is the utility that bidder i derives from not getting any non-dummy item. For each item j we are given a constant $r_j \geq 0$, called the *reserve price*, which is a lower bound on p_j. Finally, for each bidder-item pair (i, j) we are given a *utility function* $u_{i,j}(p_j)$, where p_j denotes the price of item j. The utility functions are *piece-wise linear*. That is, each $u_{i,j}(\cdot)$ is composed of $t_{i,j}$ constant-slope intervals

$$u_{i,j}^{(t)}(p_j) = v_{i,j}^{(t)} - c_{i,j}^{(t)} \cdot p_j \text{ for } p_j \in [s_{i,j}^{(t)}, e_{i,j}^{(t)}), \tag{1.6}$$

where $t \in \{1, \ldots, t_{i,j}\}$, $s_{i,j}^{(1)} = r_j$, $e_{i,j}^{(t_{i,j})} = \infty$, $s_{i,j}^{(t)} < e_{i,j}^{(t)}$ $(\forall t)$, and $e_{i,j}^{(t)} = s_{i,j}^{(t+1)}$ $(\forall t \neq t_{i,j})$. Where possible we omit (t) to improve readability. We make the following assumptions regarding the utility functions: (1) They are strictly monotonically decreasing. (2) They need *not* be globally continuous. (3) For every bidder-item pair (i, j) there exists a threshold value $\bar{p}_{i,j}$ such that $u_{i,j}(\bar{p}_{i,j}) \leq o_i$. (4) The utility functions $u_{i,j_0}(\cdot)$ for the dummy item j_0 are of the form $u_{i,j_0}(p_{j_0}) = o_i - p_{j_0}$ for $p_{j_0} \in [0, \infty)$ and $r_{j_0} = 0$.[5]

Our goal is to compute a bidder optimal solution. A solution (μ, p) consists of (1) a *matching* μ, i.e., a subset $\mu \subseteq I \times J$ of the bidder-item pairs, in which (a) every bidder i appears in *exactly one* pair $(i, j) \in \mu$ and (b) every non-dummy item $j \neq j_0$ appears in *at most one* pair, and (2) per-item *prices* $p = (p_1, .., p_k)$.[6] A solution (μ, p) is *feasible* if

$$p_{j_0} = 0 \text{ and } p_j \geq r_j \text{ for all } j \neq j_0. \tag{1.7}$$

A solution is *envy free* if it is feasible and for all i and $(i, j) \in I \times J$,

$$u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,j}(p_j), \tag{1.8}$$

where $\mu(i)$ denotes the item bidder i is matched to. A solution (μ, p) is *bidder optimal* if it is envy free and for every bidder i and every envy free solution (μ', p') we have

$$u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,\mu'(i)}(p'_{\mu'(i)}). \tag{1.9}$$

We also analyze under which circumstances a mechanism that computes a bidder optimal solution is incentive compatible. A mechanism is *incentive compatible* if for every bidder i with utility functions $u_{i,j}(\cdot)$ and every two sets of utility functions $u'_{i,j}(\cdot)$ and $u''_{i,j}(\cdot)$, where $u'_{i,j}(\cdot) = u_{i,j}(\cdot)$ for i and all j and $u'_{k,j}(\cdot) = u''_{k,j}(\cdot)$ for all $k \neq i$ and all j, and corresponding

[5] This definition together with the requirement that in every feasible solution the price of the dummy item j_0 is $p_{j_0} = 0$ ensures that in every envy free solution every bidder i has utility at least o_i.

[6] Note that although we refer to μ as matching multiple bidders i can be matched to the dummy item j_0. This could be avoided by having one dummy item per bidder, but the formulation with only one dummy item has advantages in terms of computational complexity.

solutions (μ', p') and (μ'', p'') of the mechanism we have

$$u_{i,\mu'(i)}(p'_{\mu'(i)}) \geq u_{i,\mu''(i)}(p''_{\mu''(i)}). \tag{1.10}$$

Note that this definition does not involve the reserve prices r_j or outside options o_i. This makes sense because the reserve prices r_j are typically set by the seller and misreporting o_i is never beneficial to i.[7]

1.3 Mechanism

In this section we describe and analyze our *polynomial-time mechanism* for *piece-wise linear* utilities with *non-identical slopes* and *multiple discontinuities*. We begin by showing how to reduce the problem of finding a bidder optimal solution for an input with reserve prices to the problem of finding such a solution for a different input in which the reserve prices are all zero. Afterwards we prove that the bidder optimal solution has minimal prices among all envy free solutions. We then formulate the problem as a graph problem. This allows us to define strict overdemand and to prove that an envy free solution exists if and only if no set of items is strictly overdemanded using Hall's Theorem [18]. Our mechanism starts with prices all zero and iteratively raises the prices of strictly overdemanded items. To ensure envy freeness and minimality of the resulting prices it raises the prices in an envy freeness and strict overdemand preserving manner.

1.3.1 Standard Form

We say that the input is in *standard form* if $r_j = 0$ for all j. The following lemma shows that we can without loss of generality assume that the input is in standard form as for any problem instance that is not in standard form there is a linear-time reduction to an instance in standard form. This reduction is similar to the reduction for *continuous* utility functions described in [3]. The lemma also shows that a sufficient condition for a solution (μ^*, p^*) to be bidder optimal is that the prices p^* are the minimum prices at which an envy free solution exists.[8] This was already known for *continuous* utility functions (see, e.g., [10]), but it is a novel observation for *discontinuous* utility functions.

Lemma 1. *(1) If the solution (μ, p) is bidder optimal for $u'_{i,j}(p_j) = u_{i,j}(p_j + r_j)$ and $r'_j = 0$ for all bidders i and all items j, then the solution (μ, p') with $p'_j = p_j + r_j$ for all items j is bidder optimal for $u_{i,j}(p_j)$ and r_j for all bidders i and all items j. (2) If the solution (μ^*, p^*) is envy free and $p^*_j \leq p_j$ for all items j and every envy free solution (μ, p), then (μ^*, p^*) is bidder optimal.*

Proof. First we show that if the solution (μ, p) is bidder optimal for $u'_{i,j}(p_j) = u_{i,j}(p_j + r_j)$

[7] Over-reporting can only lead to a missed chance of being assigned an item and under-reporting can only lead to a utility below the true outside option.

[8] Minimality of the prices is not a necessary condition for bidder optimality because the prices of unmatched items need not be minimal.

and $r'_j = 0$ for all bidders i and all items j, then the solution (μ, p') with $p'_j = p_j + r_j$ for all items j is bidder optimal for $u_{i,j}(p_j)$ and r_j for all bidders i and all items j. The solution (μ, p') is feasible for $u_{i,j}(p_j)$ and r_j for all bidders i and all items j because $p'_{j_0} = p_{j_0} + r_{j_0} = p_{j_0} = 0$ for the dummy item j_0 and $p'_j = p_j + r_j \geq r_j$ for all other items $j \neq j_0$. It is envy free because $u_{i,\mu(i)}(p'_{\mu(i)}) = u_{i,\mu(i)}(p_{\mu(i)} + r_{\mu(i)}) = u'_{i,\mu(i)}(p_{\mu(i)}) \geq u'_{i,j}(p_j) = u_{i,j}(p_j + r_j) = u_{i,j}(p'_j)$ for all bidders i and all items j. To see that it is bidder optimal assume by contradiction that there exists a solution (μ'', p'') that is envy free for $u_{i,j}(p_j)$ and r_j for all bidders i and all items j and has $u_{i,\mu''(i)}(p''_{\mu''(i)}) \geq u_{i,\mu(i)}(p'_{\mu(i)})$ for all bidders i; with at least one of the inequalities strict. Then the solution (μ'', p''') with $p'''_j = p''_j - r_j$ for all items j is (a) feasible for $u'_{i,j}(p_j)$ and $r'_j = 0$ for all bidders i and all items j because $p'''_{j_0} = p''_{j_0} - r_{j_0} = 0$ for the dummy item j_0 and $p'''_j = p''_j - r_j \geq 0$ for all other items $j \neq j_0$ and (b) envy free for $u'_{i,j}(p_j)$ and $r'_j = 0$ for all bidders i and all items j because $u'_{i,\mu'''(i)}(p'''_{\mu''(i)}) = u'_{i,\mu''(i)}(p''_{\mu''(i)} - r_{\mu''(i)}) = u_{i,\mu''(i)}(p''_{\mu''(i)}) \geq u_{i,j}(p''_j) = u'_{i,j}(p''_j - r_j) = u'_{i,j}(p'''_j)$ for all bidders i and all items j. Hence $u'_{i,\mu''(i)}(p'''_{\mu''(i)}) = u_{i,\mu''(i)}(p''_{\mu''(i)}) \geq u_{i,\mu(i)}(p'_{\mu(i)}) = u'_{i,\mu(i)}(p_{\mu(i)})$ for all bidders i; with at least one of the inequalities strict. This contradicts the bidder optimality of (μ, p) for $u'_{i,j}(p_j)$ and $r'_j = 0$ for all bidders i and all items j.

Next we show that if the solution (μ^*, p^*) is envy free and $p^*_j \leq p_j$ for all items j and every envy free solution (μ, p), then (μ^*, p^*) is bidder optimal. By contradiction assume that there is an envy free solution (μ', p') with $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})$ for some bidder i. Since (μ^*, p^*) is envy free, we have $u_{i,\mu^*(i)}(p^*_{\mu^*(i)}) \geq u_{i,\mu'(i)}(p^*_{\mu'(i)})$. It follows that $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu'(i)}(p^*_{\mu'(i)})$, which implies $p'_{\mu'(i)} < p^*_{\mu'(i)}$. This gives a contradiction. □

1.3.2 Graph-Theoretic Formulation

Next we formulate the problem of computing an envy free solution as a graph problem. Central to this formulation is the *first choice graph* $G_p = (I \cup J, F_p)$ at prices p, which consists of one node per bidder i, one node per item j, and an edge from i to j if and only if item j gives bidder i the highest utility at the current prices. For each bidder $i \in I$ we use $F_p(i) = \{j : \exists\, (i,j) \in F_p\}$ to denote the items demanded by this bidder at prices p and for each item $j \in J$ we use $F_p(j) = \{i : \exists\, (i,j) \in F_p\}$ to denote the bidders demanding this item at prices p. Analogously, for every subset of bidders $T \subseteq I$ we use $F_p(T) = \cup_{i \in T} F_p(i)$ to denote the set of items demanded by bidders in T at prices p and for every subset of items $S \subseteq J$ we use $F_p(S) = \cup_{j \in S} F_p(j)$ to denote be the set of bidders demanding items in S at prices p. We say that a set of non-dummy items $S \subseteq J \setminus \{j_0\}$ is *strictly overdemanded* at prices p with respect to the set of bidders $T \subseteq I$ if (a) each of the bidders in T demands only items in S, i.e., $F_p(T) \subseteq S$, and (b) for every non-empty subset R of the items in S there is more demand from bidders in T than there are items, i.e., $|F_p(R) \cap T| > |R|$. We say that S is strictly overdemanded at prices p if it is strictly overdemanded at prices p with respect to some set of bidders T. Our definition of strict overdemand is *stronger* than the definition of *overdemand* [11], which only requires that the number of bidders T demanding only items in the set S is greater than the number of items in the set. It is different from the definition of *minimal overdemand* [11], which requires that the set itself but no proper subset is overdemanded. It also differs from the definition of *directional overdemand* in [4], which is identical to our definition except that it applies to the first choice graph *after* the envisioned price increase. The advantage of our definition will become clear in the next subsection.

Lemma 2. *The following statements are equivalent: (1) The solution* (μ, p) *is envy free. (2) There exists a matching* μ *in the first choice graph* G_p *at prices* p*. (3) No set of items* $S \subseteq J \setminus \{j_0\}$ *is strictly overdemanded at prices* p*.*

Proof. We begin by showing that (1) and (2) are equivalent. A solution (μ, p) is envy free if and only if (a) $p_{j_0} = 0$ and $p_j \geq 0$ for $j \neq j_0$ and (b) $u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,j}(p_j)$ for all $(i, j) \in I \times J$. Conditions (a) and (b) are in turn satisfied if and only if all edges $(i, j) \in \mu$ belong to the first choice graph G_p at prices p with $p_{j_0} = 0$ and $p_j \geq 0$ for $j \neq j_0$.

The equivalence between (2) and (3) follows from Hall's Theorem [18], which shows that there exists a matching μ in the first choice graph at prices p with $p_{j_0} = 0$ and $p_j \geq 0$ for $j \neq j_0$ if and only if $\forall T \subseteq I$: $|F_p(T)| \geq |T|$ or $j_0 \in F_p(T)$.

Next we show that (2) implies (3). For this assume that $\forall T \subseteq I$: $|F_p(T)| \geq |T|$ or $j_0 \in F_p(T)$ and, by contradiction, that there exists a set of items $S' \subseteq J \setminus \{j_0\}$ that is strictly overdemanded with respect to the set of bidders T'. Since $S' \subseteq J \setminus \{j_0\}$ is strictly overdemanded with respect to T' we have that (a) $F_p(T') \subseteq S'$ and (b) $\forall R \subseteq S' : |F_p(R) \cap T'| > |R|$. Let $T'' = F_p(S') \cap T'$. From (a) we know that $|F_p(T')| \leq |S'|$. From (b) we know that $|T''| = |F_p(S') \cap T'| > |S'|$. Since $T'' = F_p(S') \cap T' \subseteq T'$, we have $F_p(T'') \subseteq F_p(T')$ and, thus, $|F_p(T'')| \leq |F_p(T')|$. It follows that $|T''| > |F_p(T'')|$. Since $F_p(T'') \subseteq F_p(T') \subseteq S' \subseteq J \setminus \{j_0\}$, we have $j_0 \not\in F_p(T'')$. Hence for T'' neither $|F_p(T'')| \geq |T''|$ nor $j_0 \in F_p(T'')$. This gives a contradiction.

Finally we show that (3) implies (2). For this assume that no set of items $S' \subseteq J \setminus \{j_0\}$ is strictly overdemanded and, by contradiction, that there exists $T'' \subseteq I$: $|F_p(T'')| < |T''|$ and $j_0 \not\in F_p(T'')$. Consider the smallest such T'' and some $i \in T''$. For all proper subsets $T''' \subset T'' : |F_p(T''')| \geq |T'''|$. Hence all bidders in $T'' \setminus \{i\}$ can be matched to items in $F_p(T'' \setminus \{i\})$ by Hall's Theorem [18]. Let μ' be such a matching. Let $\mu'(T'' \setminus \{i\})$ denote the items matched to bidders in $T'' \setminus \{i\}$ under μ'. Compute a maximal alternating tree $\mathscr{T}$ with respect to μ' with root i. Denote the bidders and items in this tree by $T' \subseteq T''$ and $S' = F_p(T') \subseteq F_p(T'')$. It follows that (a) $j_0 \not\in S'$ because $S' = F_p(T') \subseteq F_p(T'') \subseteq J \setminus \{j_0\}$ and (b) all items in $F_p(T'') \supseteq F_p(T') = S'$ are matched because otherwise $|F_p(T'')| > |\mu'(T'')| = |\mu'(T'' \setminus \{i\})| = |T'' \setminus \{i\}| = |T''| - 1$, i.e., $|F_p(T'')| \geq |T''|$. Hence Lemma 3 shows that S' is strictly overdemanded with respect to T'. This gives a contradiction. □

1.3.3 Alternating Paths and Trees

To identify strictly overdemanded items our mechanism makes use of alternating paths and trees: Let μ be a *partial matching*. That is, a matching in which not all of the bidders have to be matched. An *alternating path* $\mathscr{P}$ with respect to μ in the first choice graph G_p at prices p from an unmatched bidder i_0 to some item or bidder j is a sequence of edges that alternates between unmatched and matched edges and in which all items except j are non-dummy items. An *alternating tree* $\mathscr{T}$ with respect to μ with root i_0 is a tree in the first choice graph G_p at prices p which is rooted at an unmatched bidder i_0 and in which all paths from the root i_0 to a leaf j are alternating. An alternating tree is *maximal* if the first choice items of all bidders in the tree are contained in the tree and all matched items in the tree are matched to bidders in the tree. Formally: If $T \subseteq I$ and $S \subseteq J$ are the bidders and items in the tree $\mathscr{T}$, then $F_p(T) \subseteq S$ and $\mu(S) \subseteq T$. The fact that a partial matching can be augmented along an alternating path from an unmatched bidder to an unmatched item has been used before (see, e.g., [11]). The new insight of the following lemma is that there is a close correspondence between maximal alternating trees and our definition of strict overdemand.

Lemma 3. *For any maximal alternating tree $\mathscr{T}$ with respect to μ with root i_0 in G_p: (1) If the dummy item j_0 or some unmatched item $j \neq j_0$ is contained in $\mathscr{T}$, then the matching μ can be augmented along an alternating path $\mathscr{P}$ from i_0 to j_0 resp. j. (2) If all items S in $\mathscr{T}$ are non-dummy items and matched, then S is strictly overdemanded with respect to the bidders T in the tree and $|T| = |S| + 1$.*

Proof. We first show that if the dummy item j_0 or some unmatched item $j \neq j_0$ is contained in $\mathscr{T}$, then the matching μ can be augmented along an alternating path $\mathscr{P}$ from i_0 to j_0 resp. j. The path P is the path in the maximal alternating tree $\mathscr{T}$ that leads from i_0 to j_0 resp. j. All bidders on this path except i_0 are incident to two edges, one matched and one unmatched, and they are indifferent between the two. Hence we can swap the matched and unmatched edges along P to augment the size of the matching by one.

Next we show that if all items S in $\mathscr{T}$ are non-dummy items and matched, then (a) S is strictly overdemanded with respect to the bidders T in the tree and (b) $|T| = |S| + 1$. For (a) we argue as follows: We know that $j_0 \not\in S$. From the maximality of the tree $\mathscr{T}$ we get $F_p(T) \subseteq S$. We still have to show that for all $R \subseteq S : |F_p(R) \cap T| > |R|$. For every item set $R \subseteq S$ we know that there exists a node $x \in R$ such that no other node of R lies on the path P from x to the root i_0. Note that x is not the root because the root does not belong to R. Let y be the neighbor of x on P. Then y belongs to $F_p(R) \cap T$, but it is not matched to any node in R. Thus, counting the nodes matched to nodes in R and y, there are at least $|R| + 1$ nodes in $F_p(R) \cap T$. For (b) we argue as follows: By maximality of the alternating tree the items that the bidders in T are matched to must be contained in S. That all items in S are non-dummy items and matched implies that the only unmatched bidder in T is i_0. Hence $|T| = |T \setminus i_0| + 1 = |S| + 1$. □

1.3.4 Envy Freeness Preserving Price Increases

Once we have identified a strictly overdemanded set of items we need to determine how to increase the prices of the items in the set: A *price increase* d is a k-dimensional vector with entries d_j for $j \in \{1,\dots,k\}$. A price increase *preserves envy freeness* with respect to a set of first choice edges $E \subseteq F_p \setminus (I \times \{j_0\})$ at prices p if it satisfies the following conditions:

(a) The entry d_j is strictly larger than zero for items j that are the end point of at least one edge in E and it is zero otherwise.
(b) At prices $p + \lambda \cdot d$, for small enough $\lambda > 0$, each bidder i that is the end point of at least one edge in E weakly prefers each item j to which he has an edge in E to each item k that is currently among his first choice items.

Note that the dummy item j_0 is *not* the end point of an edge in E and so (a) implies that $d_{j_0} = 0$. Also note that it is sufficient to require (b) for all items that are currently among the first choice items of bidder i as bidder i will prefer each item j to which he has an edge in E to every item that was *not* among his first choice items at prices $p + \lambda \cdot d$ as long as $\lambda > 0$ is small enough. Our definition of an envy freeness preserving price increase is similar to the definition of a *competitive direction* for continuous utility functions in [3]. The next two lemmata are proved in [3] for competitive directions and continuous utility functions, we generalize them to envy freeness preserving price increases and discontinuous utility functions. The first lemma is an immediate consequence of the definition of price increases that preserve envy freeness. The second lemma gives a sufficient and necessary condition for a price increase $d \neq 0$ to preserve envy freeness for a set of first choice edges $E \subseteq F_p \setminus (I \times \{j_0\})$. It shows that a price increase d preserves envy freeness for a first choice edge $(i, j) \in F_p$ with $j \neq j_0$ if and only if the "utility drop" $c_{i,j} \cdot d_j$ on this edge is minimal across the first choice edges $(i, k) \in F_p$ incident to i. We exploit this characterization in the computation of price increases described in the next subsection.

Lemma 4. *If d is an envy freeness preserving price increase with respect to the set of first choice edges $E \subseteq F_p \setminus (I \times \{j_0\})$ at prices p, then E belongs to the set of first choice edges at prices $p + \lambda \cdot d$ for all sufficiently small $\lambda > 0$.*

Proof. Consider an arbitrary bidder-item pair $(i, j) \in E$. If $(i, k) \in F_p$: Since d preserves envy freeness for E and $(i, j) \in E$, we have $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ for all $\lambda > 0$ sufficiently small. If $(i, k) \in (I \times J) \setminus F_p$: Since $u_{i,j}(p_j) > u_{i,k}(p_k)$, we have $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ for all $\lambda > 0$ sufficiently small.[9] We conclude that $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ for all $(i, k) \in I \times J$ and all $\lambda > 0$ sufficiently small. Since $d_j \geq 0$ for all j s.t. there is a bidder i with $(i, j) \in E$ and $d_j = 0$ otherwise, we have $p_{j_0} + \lambda \cdot d_{j_0} = p_{j_0} = 0$ and $p_j + \lambda \cdot p_j \geq 0$ for all $j \neq j_0$. □

Lemma 5. *A price increase $d \neq 0$ preserves envy freeness for the set of first choice edges $E \subseteq F_p \setminus (I \times \{j_0\})$ at prices p if and only if $c_{i,j} \cdot d_j \leq c_{i,k} \cdot d_k$ for all $(i, j) \in E$ and all $(i, k) \in F_p$, where $-c_{i,j}$ and $-c_{i,k}$ are the slopes of the utility functions $u_{i,j}(\cdot)$ and $u_{i,k}(\cdot)$ at prices p.*

[9] There is no discontinuity in the utility function $u_{i,j}(\cdot)$ within the range $[p_j, p_j + \lambda \cdot d_j]$ for all $\lambda > 0$ sufficiently small because the utility function $u_{i,j}(\cdot)$ is locally right-continuous.

Proof. For the *if*-part assume that $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ for all sufficiently small $\lambda > 0$, all $(i, j) \in E$, and all $(i, k) \in F_p$. Consider arbitrary edges $(i, j) \in E$ and $(i, k) \in F_p$. By piece-wise linearity, $u_{i,j}(p_j + \lambda \cdot d_j) = u_{i,j}(p_j) - c_{i,j} \cdot \lambda \cdot d_j$ and $u_{i,k}(p_k + \lambda \cdot d_k) = u_{i,k}(p_k) - c_{i,k} \cdot \lambda \cdot d_k$ for all sufficiently small $\lambda > 0$. From this and the fact that $u_{i,j}(p_j + \lambda \cdot d_j) \geq u_{i,k}(p_k + \lambda \cdot d_k)$ it follows that $u_{i,j}(p_j) - c_{i,j} \cdot \lambda \cdot d_j \geq u_{i,k}(p_k) - c_{i,k} \cdot \lambda \cdot d_k$. Since $(i, j) \in E$ and $(i, k) \in F_p$, we have $u_{i,j}(p_j) = u_{i,k}(p_k)$ and, thus, $c_{i,j} \cdot d_j < c_{i,k} \cdot d_k$.

For the *only if*-part assume that $c_{i,j} \cdot d_j \leq c_{i,k} \cdot d_k$ for all $(i, j) \in E$ and all $(i, k) \in F_p$. Consider arbitrary edges $(i, j) \in E$ and $(i, k) \in F_p$. By piece-wise linearity, $u_{i,j}(p_j + \lambda \cdot d_j) = u_{i,j}(p_j) - c_{i,j} \cdot \lambda \cdot d_j$ and $u_{i,k}(p_k + \lambda \cdot d_k) = u_{i,k}(p_k) - c_{i,k} \cdot \lambda \cdot d_k$ for all sufficiently small $\lambda > 0$.[10] Since $(i, j) \in E$ and $(i, k) \in F_p$, we have $u_{i,j}(p_j) = u_{i,k}(p_k)$. It follows that $u_{i,j}(p_j + \lambda \cdot d_j) = u_{i,j}(p_j) - c_{i,j} \cdot \lambda \cdot d_j \geq u_{i,k}(p_k) - c_{i,k} \cdot \lambda \cdot d_k = u_{i,k}(p_k + \lambda \cdot d_k)$ for all sufficiently small $\lambda > 0$. □

1.3.5 Strict Overdemand Preserving Price Increases

It is not difficult to see that envy freeness preserving price increases are not enough to guarantee minimality of the prices. To achieve this goal we define a stronger notion of price increases, which exploits the correspondence between maximal alternating trees and strict overdemand: A *strict overdemand preserving price increase* d for a maximal alternating tree $\mathcal{T}$ with respect to μ with root i_0 in G_p with item set $S \subseteq J \setminus \{j_0\}$ and bidder set T in which all items are matched, is a price increase d such that

(a) there is some partial matching μ' that matches the same bidders and items as μ and that is identical to μ on $I \setminus T \times J \setminus S$,
(b) there is a maximal alternating tree $\mathcal{T}'$ with respect to μ' with root i_0 that has the same item and bidder set as $\mathcal{T}$, and
(c) d preserves envy freeness for the edges of the maximal alternating tree $\mathcal{T}'$.

We say that μ' is the matching that *corresponds* to d. Note that μ' can be different from μ on $T \times S$. The crucial and new fact is that by (b) all items in the tree, i.e., all items whose price is increased, remain strictly overdemanded for any small enough price increase.

Lemma 6. *If d is a strict overdemand preserving price increase for a maximal alternating tree $\mathcal{T}$ with respect to μ with root i_0 in G_p with item set $S \subseteq J \setminus \{j_0\}$ and bidder set T in which all items are matched, then S is strictly overdemanded with respect to T in $G_{p+\lambda \cdot d}$ for all sufficiently small $\lambda > 0$.*

Proof. Denote the partial matching and the maximal alternating tree corresponding to d by μ' and $\mathcal{T}'$. Since d preserves envy freeness for $\mathcal{T}'$, Lemma 4 shows that all edges in $\mathcal{T}'$ belong to the first choice graph $G_{p+\lambda \cdot d}$ at prices $p + \lambda \cdot d$ for all $\lambda > 0$ sufficiently small. Since $\mathcal{T}'$ is a maximal alternating tree with item set S and bidder set T in which all items are matched,

[10] For this we need that the intervals are closed on the left, but open on the right.

Lemma 3 shows that (1) the set of items S is strictly overdemanded with respect to the set of bidders T and (2) $|T| = |S| + 1$. □

The following lemma—our key lemma and main technical improvement over [3, 4]—shows that if strict overdemand preserving price increases are used, then the resulting prices will be minimal across all envy free solutions.

Lemma 7. *Let d be a strict overdemand preserving price increase for a maximal alternating tree $\mathcal{T}$ in G_p with item set S and bidder set T in which all items are matched. Let $\lambda > 0$ be the smallest scalar such that at $p + \lambda \cdot d$ (a) a bidder-item pair $(i, j) \in T \times J \setminus S$ enters $G_{p+\lambda \cdot d}$ or (b) the end of a constant-slope interval $u^{(t)}_{i,j}(\cdot)$ of the utility of a bidder-item $(i, j) \in T \times S$ is reached. Then for any envy free solution (μ'', p'') with $p'' \geq p$ we have $p'' \geq p + \lambda \cdot d$.*

Proof. For a contradiction suppose that $p''_s < p_s + \lambda \cdot d_s$ for some $s \in S$. Choose $\epsilon > 0$ such that $p''_s = p_s + (\lambda - \epsilon) \cdot d_s$. Note that $\epsilon \leq \lambda$ because $p''_s \geq p_s$. Let $A = \{j \in S \mid p''_j - p_j \leq (\lambda - \epsilon) \cdot d_j\}$ and let $B = F_{p+(\lambda-\epsilon)\cdot d}(A) \cap T$. Note that $A \neq \emptyset$ because $s \in A$.

Since d preserves envy freeness for $\mathcal{T}'$ and $p_j \leq p_j + (\lambda - \epsilon) \cdot d_j < p_j + \lambda \cdot d_j \leq e_j$ for all $j \in S$, i.e., for no $(i, j) \in T \times S$ there is a discontinuity in $u_{i,j}(\cdot)$ within the range $[p_j, p_j + (\lambda - \epsilon) \cdot d_j]$, we have $\mathcal{T}' \subseteq F_{p+(\lambda-\epsilon)\cdot d}$. Since $\mathcal{T}'$ covers all bidders in T and items in S and $A \subseteq S$ and S is strictly overdemanded with respect to T, we have that $|B| = |F_{p+(\lambda-\epsilon)\cdot d}(A) \cap T| > |A|$.

Next we show that $|A| \geq |F_{p''}(B)|$. For this it suffices to show that $F_{p''}(i) \subseteq A$ for all $i \in B$. For a contradiction assume that there exists an $i \in B$ and a $k \notin A$ with $k \in F_{p''}(i)$. It follows that

$$p''_k - p_k > (\lambda - \epsilon) \cdot d_k, \quad \text{and} \tag{1.11}$$

$$u_{i,k}(p''_k) \geq u_{i,j}(p''_j) \quad \text{for all } j. \tag{1.12}$$

But by the definition of B, and since $i \in B$, there must be a $j \in A$ such that $j \in F_{p+(\lambda-\epsilon)\cdot d}(i)$. It follows that

$$p''_j - p_j \leq (\lambda - \epsilon) \cdot d_j, \quad \text{and} \tag{1.13}$$

$$u_{i,j}(p_j + (\lambda - \epsilon) \cdot d_j) \geq u_{i,k}(p_k + (\lambda - \epsilon) \cdot d_k). \tag{1.14}$$

Using the fact that the utility functions are strictly monotonically decreasing we get

$$\begin{aligned} u_{i,j}(p''_j) &\geq u_{i,j}(p_j + (\lambda - \epsilon) \cdot d_j) && \text{(from (1.13))} \\ &\geq u_{i,k}(p_k + (\lambda - \epsilon) \cdot d_k) && \text{(from (1.14))} \\ &> u_{i,k}(p''_k). && \text{(from (1.11))} \end{aligned}$$

Since this would give a contradiction to (1.12), we must have $F_{p''}(B) \subseteq A$, i.e., $|A| \geq |F_{p''}(B)|$. It follows that $|B| > |A| \geq |F_{p''}(B)|$. But this shows that in (μ'', p'') not all bidders can be matched in an envy free manner. This gives a contradiction. □

1.3.6 Computing a Strict Overdemand Preserving Price Increase

Next we present a subroutine that computes a strict overdemand preserving price increase d and a corresponding matching μ' for a maximal alternating tree $\mathcal{T}$ with respect to μ with root i_0 in G_p with item set $S \subseteq J \setminus \{j_0\}$ and bidder set T in which all items are matched. The computation consists of three steps:

(1) The subroutine computes a matching σ between $T \setminus \{i_0\}$ and S consisting of first choice edges, which minimizes $\prod_{(i,j)\in\mu} c_{i,j}$, or equivalently, $\sum_{(i,j)\in\mu} \log(c_{i,j})$. It also computes an envy freeness preserving price increase d for σ. This can be accomplished by solving a linear program (LP) and its dual (DP). The duality between slopes and utility drops exploited here is reminiscent of the duality between matchings that maximize social welfare and envy free prices in [23].

(2) The subroutine extends d to an envy freeness preserving price increase for a maximal alternating tree $\mathcal{T}'$ with respect to σ with root i_0 in G_p with bidder set T and item set S.

(3) The subroutine extends σ to μ' by adding to it the bidder-item pairs from $I \setminus T \times J \setminus S$ that were matched in μ.

While (1) is essentially an application of Lemma 5 (and has been used in a similar form in [3, 4]), (2) and (3) exploit the newly established correspondence between maximal alternating trees and strict overdemand.

Subroutine for Computing a Strict Overdemand Preserving Price Increase

Input: Maximal alternating tree $\mathcal{T}$ with respect to μ with root i_0 in G_p with item set S and bidder set T in which all items are matched

Output: Strict overdemand preserving price increase d for $\mathcal{T}$ with corresponding matching μ'

Compute x as optimal solution to the following LP and let $\sigma = \{(i,j) \in T \setminus \{i_0\} \times S \mid x_{i,j} = 1\}$

$$\begin{aligned}
\min\ & \textstyle\sum_{i,j} x_{i,j} \cdot \log(c_{i,j}) \\
\text{sb}\ & \textstyle\sum_{j\in F_p(i)} x_{i,j} = 1 \quad (\forall\, i \in T \setminus \{i_0\}) \\
& \textstyle\sum_{i\in F_p(j)} x_{i,j} = 1 \quad (\forall\, j \in S) \\
& x_{i,j} \geq 0 \quad (\forall (i,j) \in F_p \cap (T \setminus \{i_0\} \times S))
\end{aligned}$$

Compute ω, ρ as optimal solution to the following DP

$$\begin{aligned}
\max\ & \textstyle\sum_i \omega_i + \sum_j \rho_j \\
\text{sb}\ & \omega_i + \rho_j \leq \log(c_{i,j}) \quad (\forall (i,j) \in F_p \cap (T \setminus \{i_0\} \times S))
\end{aligned}$$

Extend ω from $T \setminus \{i_0\}$ to T by setting $\omega_{i_0} = \min_{j\in S} \log(c_{i_0,j}) - \rho_j$

Let $H_\rho = (T \cup S, E_\rho)$, where $E_\rho = \{(i,j) \in F_p \cap (T \times S) \mid \omega_i + \rho_j = \log(c_{i,j})\}$

5 Let $\mathscr{T}'$ be a maximal alternating tree in H_ρ with respect to σ with root i_0
6 Let $S' \subseteq S$ and $T' \subseteq T$ denote the items and bidders in $\mathscr{T}'$
7 **while** $T' \neq T$ or $S' \neq S$ **do**
8 Let $\delta = \min_{(i,j)\in F_p: i\in T', j\in S\setminus S'} \log(c_{i,j}) - \omega_i - \rho_j$
9 Set $\rho_j = \rho_j + \delta$ for all $j \in S \setminus S'$ and $\omega_i = \omega_i - \delta$ for all $i \in T \setminus T'$
10 Recompute $\mathscr{T}'$, T', and S'
11 **end while**
12 Set $d_j = e^{-\rho_j}$ for all $j \in S$ and $d_j = 0$ otherwise
13 Set $\mu' = \sigma \cup (\mu \cap (I \setminus T \times J \setminus S))$
14 Output d and μ'

Lemma 8. *This subroutine finds a strict overdemand preserving price increase and a corresponding matching. It can be implemented to run in time* $O(\min(n,k)^3)$.

Proof. Let x, ω, and ρ be defined as in the mechanism. The constraint matrix of LP is totally unimodular, i.e., $x_{i,j} \in \{0,1\}$ for all $(i,j) \in F_p \cap (T \setminus \{i_0\} \times S)$ [22]. Hence $\sum_{j\in F_p(i)} x_{i,j} = 1$ $(\forall i \in T\setminus\{i_0\})$ and $\sum_{i\in F_p(j)} x_{i,j} = 1$ $(\forall j \in S)$ ensure that $\sigma = \{(i,j) \in T\setminus\{i_0\} \times S \mid x_{i,j} = 1\}$ matches every bidder $i \in T \setminus \{i_0\}$ and every item $j \in S$ exactly once. From duality:

1. For all $i \in T \setminus \{i_0\}$ and all $j \in S$ with $(i,j) \in \sigma$: $\omega_i + \rho_j = \log(c_{i,j})$.
2. For all $i \in T \setminus \{i_0\}$ and all $j \in S$: $\omega_i + \rho_j \leq \log(c_{i,j})$.

If we extend ω from $T \setminus \{i_0\}$ to T as described in l. 3, then we also have:

3. There exists a $j \in S$: $\omega_{i_0} + \rho_j = \log(c_{i_0,j})$.
4. For all $j \in S$: $\omega_{i_0} + \rho_j \leq \log(c_{i_0,j})$.

Let $H_\rho = (S \cup T, E_\rho)$, E_ρ, $\mathscr{T}'$, T' and S' be defined as in ll. 4-6. It is not difficult to see that the while-loop in ll. 7-11 has the following properties:

a. For all $i \in T'$ and all $j \in S'$: If we had $\omega_i + \rho_j <$ (resp. $=$) $\log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j <$ (resp. $=$) $\log(c_{i,j})$ after the update.
b. For all $i \in T \setminus T'$ and all $j \in S \setminus S'$: If we had $\omega_i + \rho_j <$ (resp. $=$) $\log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j <$ (resp. $=$) $\log(c_{i,j})$ after the update.
c. For all $i \in T \setminus T'$ and all $j \in S'$: If we had $\omega_i + \rho_j \leq \log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j < \log(c_{i,j})$ after the update.
d. For all $i \in T'$ and $j \in S \setminus S'$: If we had $\omega_i + \rho_j < \log(c_{i,j})$ before the update, then we have $\omega_i + \rho_j \leq \log(c_{i,j})$ after the update.
e. There exist at least one $i \in T'$ and $j \in S \setminus S'$: $\omega_i + \rho_j = \log(c_{i,j})$ after the update.

From a. to e. we get that no edge from σ and $\mathscr{T}'$ in E_ρ is lost. From e. we get that at least one edge from some $i \in T'$ to some $j \in S \setminus S'$ is added to E_ρ. Since this item j was matched under σ

along an edge in E_ρ to an item $i' \in T \setminus T'$, we know that after each iteration of the while-loop the maximal alternating tree $\mathcal{T}'$ with respect to σ with root i_0 in H_ρ will at least cover the bidders and items in $T' \cup \{i'\}$ and $S' \cup \{j\}$. Hence, after at most $|S| \le k$ iterations, $\mathcal{T}'$ will cover the same bidders and items as $\mathcal{T}$.

Let d be defined as in l. 12. Then $d \neq 0$ preserves envy freeness for $\mathcal{T}'$ by Lemma 5 because:

1. For all $(i,j) \in \mathcal{T}'$ and all $(i,k) \in F_p$ we have that $c_{i,j} \cdot d_j \le c_{i,k} \cdot d_k$ because $e^{\omega_i} = e^{\log(c_{i,j})-\rho_j} = c_{i,j} \cdot e^{-\rho_j} = c_{i,j} \cdot d_j$ and $e^{\omega_i} \le e^{\log(c_{i,k})-\rho_k} = c_{i,k} \cdot e^{-\rho_k} = c_{i,k} \cdot d_k$.
2. For all $(i,j) \in \mu' \setminus \mathcal{T}'$ and $(i,k) \in F_p$ we trivially have that $c_{i,j} \cdot d_j \le c_{i,k} \cdot d_k$ because $d_j = 0$, $d_k \ge 0$, and $c_{i,k} \ge 0$.

Let μ' be defined as in l. 13, then μ' matches the same bidders and items as μ because (1) μ' is identical to σ on $T \times S$ and (2) μ' is identical to μ on $I \setminus T \times J \setminus S$.

The LP and the DP can be solved in time $O(\min(n,k)^3)$ [16, 24]. The maximal alternating tree $\mathcal{T}'$ can be computed in time $O(\min(n,k)^2)$ using a breadth-first search approach. The while-loop in ll. 7-11 can be implemented using "slack variables" $\delta_j = \min_{(i,j)\in F_p: i\in T'}(\log(c_{i,j}) - \omega_i - \rho_j)$ for each item $j \in S \setminus S'$ so that *all* iterations of the while loop take total time $O(\min(n,k)^2)$: The initialization of the δ_j's takes time $O(\min(n,k)^2)$ as for each of the up to $\min(n,k)$ items in $S \setminus S'$ the minimum is computed over the up to $\min(n,k)$ bidders in T'. In each iteration of the while-loop at least one bidder-item pair $(i,j) \in T' \times S \setminus S'$ is added to $T' \times S'$. Since $|S \setminus S'| \le \min(n,k)$ it follows that there are at most $\min(n,k)$ iterations. Using the δ_j's for $j \in S \setminus S'$ the δ in l.8 can be computed in time $O(\min(n,k))$. When the ω_i's and ρ_j's are updated in l. 9, the δ_j's are adapted in time $O(\min(n,k))$ by subtracting δ from each δ_j. Thus, ll. 8 and 9 take time $O(\min(n,k))$ per iteration for a total of $O(\min(n,k)^2)$. Instead of re-computing the maximal alternating tree $\mathcal{T}'$ in l. 10 from scratch we can keep the old one and add the required edges. Thus maintaining $\mathcal{T}'$ takes only time $O(\min(n,k)^2)$ for all iterations of the while loop. Additionally the δ_j's must be updated. For each bidder that is added to T' all δ_j's must be updated. This takes time $O(\min(n,k))$ per bidder. But this happens at most once for each of the up to $\min(n,k)$ bidders that are added to T', since no bidder is ever removed from T'. Thus, all the updates to the δ_j's that are required when bidders are added to T' take total time $O(\min(n,k)^2)$. □

1.3.7 Computing a Bidder Optimal Solution

Our mechanism starts with an empty matching $\mu = \emptyset$ and prices $p = 0$. It then matches one bidder after the other until eventually all bidders are matched. For this it computes a maximal alternating tree $\mathcal{T}$ with respect to μ with root i_0, where i_0 is the bidder to be matched, in the first choice graph G_p. If the alternating tree contains the dummy item j_0 or an unmatched item j, then by Lemma 3 the current matching μ can be augmented along an alternating path from i_0 to j_0 resp. j. If this is not the case, then—again by Lemma 3—the items S in the tree are strictly overdemanded with respect to the bidders T in the tree. In this case the mechanism computes a strict overdemand preserving price increase d together with a corresponding matching μ' and raises the prices in compliance with d until (a) a bidder-item pair $(i,j) \in T \times J \setminus S$ enters the first choice graph $G_{p+\lambda\cdot d}$ or (b) the end of a constant-slope interval $u_{i,j}^{(t)}(\cdot)$ of the utility of a bidder-item pair $(i,j) \in T \times S$ is reached. In either case the current matching μ is replaced with μ' and the minimality of the new prices is guaranteed by Lemma 7. If at least one of the new prices $p_j + \lambda \cdot d_j$ corresponds to a discontinuity, then one or multiple edges might drop out of the first choice graph. The mechanism corrects for this by removing such edges from the matching if necessary. If no discontinuity is reached, then the maximal alternating tree $\mathcal{T}$ rooted at i_0 grows by at least one item.

Mechanism for Computing a Bidder Optimal Solution

Input: Bidders I, items J, piece-wise linear utilities $u_{i,j}(\cdot)$ with non-identical slopes and multiple discontinuities, reserve prices $r_j = 0$, outside options o_i
Output: Bidder optimal solution (μ, p)

Set $p_j = 0$ for all j and set $\mu = \emptyset$
while There exists an unmatched bidder i_0 **do**
  Compute maximal alternating tree $\mathcal{T}$ wrt μ in the first choice graph G_p with root i_0
  Let T and S be the bidders and items in $\mathcal{T}$
  while All items in S are matched and S does not contain the dummy item j_0 **do**
    Compute a strict overdemand preserving price increase d for $\mathcal{T}$ and a corresponding matching μ' (using the subroutine from the previous subsection)
    Let $\lambda > 0$ be the smallest scalar such that at prices $p + \lambda \cdot d$
      (a) a bidder-item pair $(i,j) \in T \times J \setminus S$ enters the first choice graph $G_{p+\lambda\cdot d}$ or
      (b) the end of a constant-slope interval $u_{i,j}^{(t)}(\cdot)$ of the utility of a bidder-item pair $(i,j) \in T \times S$ is reached
    Set $p_j = p_j + \lambda \cdot d_j$ for all $j \in J$ and set $\mu = \mu'$
    Remove bidder-item pairs from μ that do not belong to the first choice graph G_p
    Compute maximal alternating tree $\mathcal{T}$ wrt μ in the first choice graph G_p with root i
    Let T and S be the bidders and items in $\mathcal{T}$
  end while
  Augment μ along alternating path P from i_0 to unmatched item j or dummy item j_0
end while
Output (μ, p)

Theorem 1. *This mechanism finds a bidder optimal solution. It can be implemented to run in time* $O((n \cdot \min(n,k) + D \cdot \min(n,k) + T) \cdot \min(n,k) \cdot (\min(n,k)^2 + k))$, *where* $D = \sum_{i,j} d_{i,j}$ *and* $T = \sum_{i,j} t_{i,j}$ *denote the total number of discontinuities and constant-slope intervals.*

Proof. The matching μ is a subset of the first choice edges at prices p. Hence Lemma 2 shows that (μ, p) is envy free. By Lemma 1 (μ, p) is bidder optimal if $p_j \le p''_j$ for every item j and every envy free solution (μ'', p''). Let $p^{(t)}$ denote the prices after the t-th update. We prove that

$p_j^{(t)} \le p''_j$ for every item j, every envy free solution (μ'', p''), and all time steps t by induction over t.

Induction basis: For $t = 0$ the claim follows from the fact that every envy free solution (μ'', p'') has $p''_j \ge 0$ for all j.

Inductive step: To prove the claim for t assume that it holds for $t-1$. Let $\mathcal{T}$ be the maximal alternating tree with respect to the current matching $\mu^{(t-1)}$ with root i_0 right before the t-th update. Let S and T denote the items and bidders in $\mathcal{T}$. Let d be a strict overdemand preserving price increase for $\mathcal{T}$ with corresponding alternating tree $\mathcal{T}'$ and matching μ'. Let λ be defined as in the mechanism. Note that λ can be computed in time $O(\min(n,k)\cdot k)$ by iterating over all bidders in T, of which there are at most $\min(n,k)$, and all items in J. The mechanism sets $p_j^{(t+1)} = p_j^{(t)}$ for $j \notin S$ and $p_j^{(t+1)} = p_j^{(t)} + \lambda \cdot d_j$ for $j \in S$. Lemma 7 shows that any envy free solution (μ'', p'') with $p'' \ge p^{(t)}$ must have $p''_j \ge p^{(t)} + \lambda \cdot d_j$. It follows that $p''_j \ge p_j^{(t+1)}$ for all j.

We bound the total time required by (1) the outer while-loop *without* the inner while-loop (ll. 2-4 & 13-14) and the inner while-loop corresponding to Case (a) (ll. 5-12) separately from the total time required by (2) the inner while-loop corresponding to Case (b) (ll. 5-12).

To obtain a bound for (1) observe that: (i) In each iteration of the outer while-loop exactly one bidder gets matched. (ii) Bidders can get unmatched only if the boundary of a box is reached that corresponds to a discontinuity in at least one of the utility functions $u_{i,j}(\cdot)$. (iii) A discontinuity in $u_{i,j}(\cdot)$ can only unmatch bidder i. (iv) Since the prices are monotonically increasing at most $O(D)$ discontinuities are reached. From (i) to (iv) we deduce that there are at most $O(n+D)$ iterations of the outer while-loop without the inner while-loop. Each iteration of the outer while-loop without the inner while-loop takes time $O(\min(n,k)^2)$. From (i) to (iv) it also follows that there are at most $O(n+D)$ iterations of the inner while corresponding to Case (a) such that right before their execution either the outer while-loop was executed or the inner while-loop was executed and a discontinuity was reached. Between any two such iterations there can be at most $O(\min(n,k))$ iterations of the inner while-loop corresponding to Case (a) because each of these iterations adds at least one item to the maximal alternating tree under consideration. Each iteration of the inner while-loop takes time $O(\min(n,k)\cdot(\min(n,k)^2+k))$, namely $O(\min(n,k)^3)$ for computing the overdemand preserving price increase (see Lemma 8) and $O(\min(n,k)\cdot k)$ for computing the λ value. Hence a bound for (1) is $O((n\cdot \min(n,k) + D\cdot \min(n,k))\cdot \min(n,k)\cdot(\min(n,k)^2+k))$.

To obtain a bound for (2) observe that because the prices are monotonically increasing there are at most $O(T)$ iterations of the inner-while loop that correspond to Case (b). As argued above each iteration of the inner while-loop takes time $O(\min(n,k)\cdot(\min(n,k)^2+k))$. Hence a bound for (2) is $O(T\cdot \min(n,k)\cdot(\min(n,k)^2+k))$. □

1.4 Incentive Compatibility

In this section we analyze under which conditions mechanisms that compute a bidder optimal solution for piece-wise linear utilities with non-identical slopes and multiple discontinuities are *incentive compatible*. We begin by providing an example that shows that for unrestricted inputs bidder optimality does *not* imply incentive compatibility. In this example, if our mechanism is used to compute a bidder optimal solution, two discontinuities are reached in the same price increase. Motivated by this observation we develop a condition on the input, which we refer to as *general position*, that ensures that at most one discontinuity is reached in each price increase of our mechanism. Afterwards, we show that this is enough to guarantee the existence of a bidder optimal solution with a specific structure. Finally, we use the existence of this specific bidder optimal solution to show that for inputs in general position bidder optimality implies incentive compatibility.

1.4.1 Counterexample

The following example shows that in the presence of discontinuities our mechanism, although it finds a bidder optimal solution, need not be incentive compatible. It even shows that this is true for any mechanism that computes a bidder optimal solution as the bidder optimal utilities for both the truthful input as well as the falsified input are unique.[11]

Example 3. Lying pays off (This example shows that bidder optimality does not imply incentive compatibility.) There are two bidders $i \in \{1,2\}$ and two items $j \in \{1,2\}$. The utility functions for $i \in \{1,2\}$ are:

$$u_{i,1}(p_1) = \begin{cases} 20 - p_1 & \text{for } p_1 \in [0,5), \text{and} \\ -\infty & \text{otherwise,} \end{cases}$$
$$u_{i,2}(p_2) = \quad 1 - p_2 \quad \text{for } p_2 \in [0,\infty).$$

The reserve prices are $r_j = 0$ for $j \in \{1,2\}$ and the outside options are $o_i = 0$ for $i \in \{1,2\}$.

Given this input our mechanism computes a bidder optimal solution as follows: Starting from prices all zero, at which both bidders strictly prefer item 1 over item 2, it raises the price of the first item to 5. At this point both bidders lose their interest in item 1 and demand item 2 instead. The mechanism continues by raising the price of item 2 to 1. Now both bidders are indifferent between item 2 and the dummy item. Hence the mechanism can assign bidder 1 the dummy item and bidder 2 item 2 (or vice versa). For this solution both bidders have a utility of zero.

Next suppose bidder 1 lies about his utility function for item 1 by reporting $u_{1,1} = 0 - p_1$ for $p_1 \in [0,\infty)$. In this case bidder 1 (presumably) prefers item 2 over item 1 and bidder 2 prefers

[11] Note that although the bidder optimal utilities are unique the bidder optimal solutions are *not*. For the truthful input, for example, at the prices computed by our mechanism we can either match bidder 1 to item 2 and leave bidder 2 unmatched or vice versa.

item 1 over item 2 when prices are all zero. Hence our mechanism will leave the price of all items at 0 and assign item 1 to bidder 2 and item 2 to bidder 1. Bidder 1's utility for this solution is 1 and, thus, strictly higher than the utility that he would get if he reported truthfully.

The crucial point—as we will show below—is that in the computation of the bidder optimal solution for the original utility functions the first choice edges from bidder 1 to item 1 and from bidder 2 to item 1 broke away in the same price increase because the corresponding discontinuities were reached.

1.4.2 Price-Independent Formulation

We define next a condition on the input that ensures that in each price increase of our mechanism at most one edge breaks away due to a discontinuity. Which edges break away depends on the current prices and the price increases. However, using the following idea we can write down a condition that does *not* depend on the current prices: Suppose that the edges (i, j), (i', j), and (i', j') belong to the first choice graph G_p at prices p. Then,

$$v_{i',j} - c_{i',j} \cdot p_j = v_{i',j'} - c_{i',j'} \cdot p_{j'}. \tag{1.15}$$

Suppose further that d is an envy free price increase for the set of first choice edges $E = \{(i, j), (i', j), (i', j')\}$. Then,

$$v_{i',j} - c_{i',j} \cdot (p_j + \lambda d_j) = v_{i',j'} - c_{i',j'} \cdot (p_{j'} + \lambda d_{j'}). \tag{1.16}$$

By subtracting (1.15) from (1.16), dividing by $\lambda > 0$, and after rearranging we get

$$c_{i',j} \cdot d_j = c_{i',j'} \cdot d_{j'}. \tag{1.17}$$

Now suppose that the discontinuities $D_{i,j}$ and $D_{i',j'}$ are reached simultaneously. Then by (1.16):

$$v_{i',j} - c_{i',j} \cdot D_{i,j} = v_{i',j'} - c_{i',j'} \cdot D_{i',j'}. \tag{1.18}$$

Equation (1.17) lets us divide the left hand side of this equation by $c_{i',j} \cdot d_j$ and the right hand side by $c_{i',j'} \cdot d_{j'}$. Subtracting $(1/d_j) \cdot (v_{i,j}/c_{i,j})$ from both sides of the resulting equation and rearranging terms gives

$$\frac{1}{d_j}\left(D_{i,j} - \frac{v_{i,j}}{c_{i,j}}\right) = -\frac{1}{d_j}\frac{v_{i,j}}{c_{i,j}} + \frac{1}{d_j}\frac{v_{i',j}}{c_{i',j}} + \frac{1}{d_{j'}}\left(D_{i',j'} - \frac{v_{i',j'}}{c_{i',j'}}\right).$$

Below we will define a weighted multigraph and alternating walks in this multigraph such that the left and right hand side of the preceding equation correspond to the weights of two alternating walks in this multigraph, namely $P = (i, j)$ and $Q = (i, j, i', j')$. Note that neither the weight of P nor the weight of Q depends on the prices.

1.4.3 General Position

For a given input we define a multigraph, called *input graph*, as follows: There is one node per bidder i and one node per item j. There are three types of edges: (1) There is a *forward edge* from bidder i to item j for each constant-slope interval of $u_{i,j}(\cdot)$. (2) There is a *backward edge* from item j to bidder i for each constant-slope interval of $u_{i,j}(\cdot)$. (3) There is a *discontinuity edge* from bidder i to item j for each discontinuity $D_{i,j}$ of $u_{i,j}(\cdot)$.

Let $P = (i_0, j_1, \ldots, i_s, j_s)$ be a walk in the input graph that alternates between forward and backward edges, and ends with a discontinuity edge. Let d be a price increase such that $d_j = (c_{i,j'}/c_{i,j}) \cdot d_{j'}$ for any two edges (i, j) and (i, j') on P. Define the weight of each forward edge (i, j) on P with respect to d as $(-1/d_j) \cdot (v_{i,j}/c_{i,j})$, of each backward edge (j, i) as $(1/d_j) \cdot (v_{i,j}/c_{i,j})$, and of the discontinuity edge (i, j) as $(1/d_j) \cdot (D_{i,j} - v_{i,j}/c_{i,j})$. Here $v_{i,j}$ and $c_{i,j}$ are the constants of the corresponding constant-slope interval. Define the weight $w_d(P)$ of P with respect to d as the sum of these weights.

An input is in *general position* if for no two walks P and Q that start with the same bidder and end with a distinct discontinuity edge and for no price increase d such that $d_j = (c_{i,j'}/c_{i,j}) \cdot d_{j'}$ for any two edges (i, j) and (i, j') on P resp. Q we have $w_d(P) = w_d(Q)$.[12]

1.4.4 Properties of the Bidder Optimal Solution

We now show that for inputs in general position at most one discontinuity is reached in each price increase of our mechanism. Afterwards, we argue that in this case items that get unmatched can get matched again in the subsequent iteration. Finally, we show that if our mechanism is modified accordingly, then it finds a bidder optimal solution with the following properties: (i) The price p_j of all unmatched items j is $p_j = r_j$. (ii) There is at least one matched item j with $p_j = r_j$.[13]

Lemma 9. *If the input is in general position, then in each price increase of our mechanism at most one discontinuity is reached.*

Proof. For a contradiction suppose that two discontinuities are reached in the same price increase. Without loss of generality assume that they are reached on (i_1, j_1) and (i_t, j_{t-1}), and denote them by D_{i_1,j_1} and $D_{i_t,j_{t-1}}$. Consider the alternating walks $P = (i_1, j_1)$ and $Q = (i_1, j_1, i_2, j_2, \ldots, i_{t-1}, j_{t-1}, i_t, j_{t-1})$ in the maximal alternating tree corresponding to the price increase. Note that these walks always exist because (1) either j_1 lies on the path from i_1 to i_t or $i_2 = i_1$ and (2) either j_{t-1} lies on the path from i_1 to i_t or $i_t = i_{t-1}$. Since the discontinuities are reached at the same time there must be a price increase $d \neq 0$ that preserves envy freeness for the maximal alternating tree containing P and Q and a positive scalar $\lambda > 0$ such that

[12] This definition is more general than the definition of general position in [1] in that it also applies to piece-wise linear utilities with non-identical slopes and multiple discontinuities.

[13] Note that the first condition is exactly what is needed to ensure that the bidder optimal solution coincides with the competitive equilibrium that maximizes the payoffs to the bidders.

$\lambda \cdot d_{j_1} = D_{i_1,j_1} - p_{j_1}$ and $\lambda \cdot d_{j_{t-1}} = D_{i_t,j_{t-1}} - p_{j_{t-1}}$. Dividing the former by d_{j_1} and the latter by $d_{j_{t-1}}$ and equating the resulting equations gives

$$\frac{1}{d_{j_1}} \cdot D_{i_1,j_1} - \frac{1}{d_{j_{t-1}}} \cdot D_{i_t,j_{t-1}} = \frac{1}{d_{j_1}} \cdot p_{j_1} - \frac{1}{d_{j_{t-1}}} \cdot p_{j_{t-1}}. \tag{1.19}$$

From the fact that P and Q belong to the first choice graph G_p at prices p we get that

$$v_{i_{s+1},j_s} - c_{i_{s+1},j_s} \cdot p_{j_s} = v_{i_{s+1},j_{s+1}} - c_{i_{s+1},j_{s+1}} \cdot p_{j_{s+1}} \quad \text{for } s = 1..t-2. \tag{1.20}$$

Solving for p_{j_s} we get that

$$p_{j_s} = \frac{v_{i_{s+1},j_s}}{c_{i_{s+1},j_s}} - \frac{v_{i_{s+1},j_{s+1}}}{c_{i_{s+1},j_s}} + \frac{c_{i_{s+1},j_{s+1}}}{c_{i_{s+1},j_s}} \cdot p_{j_{s+1}} \quad \text{for } s = 1..t-2. \tag{1.21}$$

From the fact that P and Q also belong to the first choice graph $G_{p+\lambda\cdot d}$ at prices $p+\lambda\cdot d$ we get that $v_{i_{s+1},j_s} - c_{i_{s+1},j_s}\cdot(p_{j_s}+\lambda\cdot d_{j_s}) = v_{i_{s+1},j_{s+1}} - c_{i_{s+1},j_{s+1}}\cdot(p_{j_{s+1}}+\lambda\cdot d_{j_{s+1}})$ for $s = 1..t-2$. Subtracting (1.20) and solving for $d_{j_{s+1}}$ we get that

$$d_{j_{s+1}} = \frac{c_{i_{s+1},j_s}}{c_{i_{s+1},j_{s+1}}} \cdot d_{j_s} \quad \text{for } s = 1..t-2. \tag{1.22}$$

Solving the recurrence (1.21) for p_{j_1}, substituting (1.22), and rearranging gives:

$$\frac{1}{d_{j_1}} \cdot p_{j_1} - \frac{1}{d_{j_{t-1}}} \cdot p_{j_{t-1}} = \sum_{s=1}^{t-2} \left(\frac{1}{d_{j_{s+1}}} \cdot \frac{v_{i_{s+1},j_s}}{c_{i_{s+1},j_{s+1}}} - \frac{1}{d_{j_{s+1}}} \cdot \frac{v_{i_{s+1},j_{s+1}}}{c_{i_{s+1},j_{s+1}}} \right).$$

We combine this with (1.19) to get

$$\frac{1}{d_{j_1}} \cdot D_{i_1,j_1} - \frac{1}{d_{j_{t-1}}} \cdot D_{i_t,j_{t-1}} = \sum_{s=1}^{t-2} \left(\frac{1}{d_{j_{s+1}}} \cdot \frac{v_{i_{s+1},j_s}}{c_{i_{s+1},j_{s+1}}} - \frac{1}{d_{j_{s+1}}} \cdot \frac{v_{i_{s+1},j_{s+1}}}{c_{i_{s+1},j_{s+1}}} \right).$$

We add $(-1/d_{j_1})\cdot(v_{i_1,j_1}/c_{i_1,j_1})$ to both sides and $(1/d_{j_t})\cdot(v_{i_t,j_{t-1}}/c_{i_t,j_t}) - (1/d_{j_{t-1}})\cdot(v_{i_t,j_{t-1}}/c_{i_t,j_{t-1}}) = 0$ to the right hand side. After rearranging we get

$$\frac{1}{d_{j_1}} \cdot \left(D_{i_1,j_1} - \frac{v_{i_1,j_1}}{c_{i_1,j_1}} \right) = \sum_{s=1}^{t-1} \left(-\frac{1}{d_{j_s}} \cdot \frac{v_{i_s,j_s}}{c_{i_s,j_s}} + \frac{1}{d_{j_{s+1}}} \cdot \frac{v_{i_{s+1},j_s}}{c_{i_{s+1},j_{s+1}}} \right) + \frac{1}{d_{j_{t-1}}} \cdot \left(D_{i_t,j_{t-1}} - \frac{v_{i_t,j_{t-1}}}{c_{i_t,j_{t-1}}} \right).$$

Since the left hand side corresponds to $w_d(P)$ and the right hand side to $w_d(Q)$, we get a contradiction to the fact that the input is in general position. □

Lemma 10. *If the input is in general position and our mechanism is used, then every item that gets unmatched can get matched again in the subsequent iteration.*

Proof. Consider an arbitrary price increase in which a matched item gets unmatched. Let i_0 be the bidder and let μ be the matching under consideration. Denote the maximal alternating tree with respect to μ with root i_0 by $\mathcal{T}$. Let d be the overdemand preserving price increase for $\mathcal{T}$ computed by our mechanism. Denote the corresponding maximal alternating tree and matching by $\mathcal{T}'$ resp. μ'. We know that μ' matches the same set of bidders and items as μ. A matched item j can get unmatched only if the edge along which this item is matched under μ' drops out of the first choice graph $G_{p+\lambda\cdot d}$ for some $\lambda > 0$. Lemma 9 shows that apart from this edge no other edge can drop out of the first choice graph at this point. But then, because all items in $\mathcal{T}'$ are closer to the root i_0 than the bidder they are matched to, there must be an alternating path P from the unmatched item j to the root i_0. Hence in the subsequent iteration item j can get matched again. □

Proposition 1. *Suppose that the input is in general position and that our mechanism is modified so that items that get unmatched get matched again in the subsequent iteration. Then in the bidder optimal solution* (μ, p) *computed by our mechanism (i)* $p_j = r_j$ *for all unmatched items* j *and (ii)* $p_j = r_j$ *for at least one matched item* j.

Proof. Suppose that our mechanism is used to compute a bidder optimal solution (μ, p). By Lemma 10 we can assume that items that get unmatched get matched again in the subsequent iteration. This assumption has two implications: (a) Every item that ever got matched will be matched in the end. (b) The last item that gets matched was not matched before.

Our mechanism initializes the prices of all items to the reserve prices, and it raises the price of an item strictly above its reserve price only if the item belongs to the set of items S in a maximal alternating tree $\mathcal{T}$ in which all items are matched. Together with (a) this shows that $p_j > r_j$ implies that item j is matched or, conversely, that all unmatched items j have $p_j = r_j$. Together with (b) this shows that the last item j that gets matched has $p_j = r_j$, which proves the claim about the matched items. □

1.4.5 Characterization

Next we prove that for inputs in general position every mechanism that computes a bidder optimal solution is incentive compatible. Using the existence of a bidder optimal solution with a specific structure established in the previous subsection we first show that (a) no feasible solution can be strictly better than the bidder optimal solution for *all* bidders and (b) if a feasible solution is strictly better for *some* bidders, then there must be a bidder that is *not* better off that is *not* envy free. Afterwards, we observe that if we assume that a subset of the bidders benefits from misreporting their utilities, then the bidder optimal solution for the reported utilities is feasible for the true utilities. Hence (a) leads to a contradiction if *all* bidders benefitted from misreporting. Otherwise, if only *some* bidders benefitted from misreporting, (b) shows that there must be a bidder that is *not* better off that is *not* envy free, which contradicts the fact that the mechanism computes a bidder optimal and, a forteriori, envy free solution for the reported utilities.[14]

[14] Note that this argument actually shows the stronger claim that for inputs in general position no *group* of bidders can misreport their utilities to achieve a strictly higher utility.

Lemma 11. *Suppose that the input is in general position and that the solution* (μ^*, p^*) *is bidder optimal. Then no feasible solution* (μ', p') *can have* $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})$ *for all* i.

Proof. Since the input is in general position Proposition 1 shows the existence of a bidder optimal solution (μ, p) such that (i) $p_j = r_j$ for all unmatched items j and (ii) $p_j = r_j$ for at least one matched item j. For a contradiction assume that there is a feasible solution (μ', p') with $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})$ for all i. Since $u_{i,\mu(i)}(p_{\mu(i)}) = u_{i,\mu^*(i)}(p^*_{\mu^*(i)})$ for all i, it follows that $u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu(i)}(p_{\mu(i)})$ for all i. Consider any pair $(i, j) \in \mu'$. Then $u_{i,j}(p'_j) = u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,j}(p_j)$ and, thus, $p_j > p'_j \geq r_j$. Hence, by condition (i), item j must be matched under μ. We conclude that (a) all items that are matched under μ' are also matched under μ and (b) $p'_j < p_j$ for all of these items j.

Case 1: At least one bidder i is matched to the dummy item j_0 under μ'. By condition (b) $p_{j_0} > p'_{j_0} \geq 0$, which contradicts the feasibility of the solution (μ, p).

Case 2: All bidders i are matched to non-dummy items j under μ'. By condition (a) all bidders are matched to non-dummy items under μ. Condition (ii) shows that at least one item j is matched under μ at price $p_j = r_j$. But then condition (b) shows that $p'_j < p_j = r_j$, which contradicts the feasibility of the solution (μ', p'). □

Lemma 12. *Suppose the input is in general position, the solution* (μ^*, p^*) *is bidder optimal, the solution* (μ', p') *is feasible, and* $I^+ = \{i \in I \mid u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})\} \neq \emptyset$. *Then there exists a bidder-item pair* $(i, j) \in I \setminus I^+ \times J$ *such that* $u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j)$.

Proof. Since the input is in general position Proposition 1 shows the existence of a bidder optimal solution (μ, p) such that (i) $p_j = r_j$ for all unmatched items j and (ii) $p_j = r_j$ for at least one matched item j. Since $u_{i,\mu(i)}(p_{\mu(i)}) = u_{i,\mu^*(i)}(p^*_{\mu^*(i)})$ for all i, we have $I^+ = \{i \in I \mid u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu(i)}(p_{\mu(i)})\}$. Let $\mu(I^+)$ resp. $\mu'(I^+)$ denote the set of items matched to bidders in I^+ under μ resp. μ'. From Lemma 11 we know that $I^+ \neq I$.

Case 1: $\mu(I^+) \neq \mu'(I^+)$. There must be an item $j \in \mu'(I^+)$ such that $j \notin \mu(I^+)$. Let $i' \in I^+$ be the bidder that is matched to item j in μ'. Since $i' \in I^+$ and the solution (μ, p) is envy free we have that $u_{i',j}(p'_j) = u_{i',\mu'(i')}(p'_{\mu'(i')}) > u_{i',\mu(i')}(p_{\mu(i')}) \geq u_{i',\mu'(i')}(p_{\mu'(i')}) = u_{i',j}(p_j)$ which shows that $p_j > p'_j \geq r_j$. Thus, by condition (i), item j is matched under μ. Let $i \in I \setminus I^+$ be the bidder that is matched to item j under μ. Since $i \notin I^+$ it follows that $u_{i,\mu'(i)}(p'_{\mu'(i)}) \leq u_{i,\mu(i)}(p_{\mu(i)}) = u_{i,j}(p_j) < u_{i,j}(p'_j)$.

Case 2: $\mu(I^+) = \mu'(I^+)$. Let $J^+ = \mu(I^+) = \mu'(I^+)$. Consider the following *restricted problem*: The set of bidders is I^+, the set of items is J^+, the utility functions are $u^+_{i,j}(\cdot) = u_{i,j}(\cdot)$ for all $(i, j) \in I^+ \times J^+$, the reserve prices are $r^+_j = \max(r_j, \max_{i \notin I^+}(u^{-1}_{i,j}(u_{i,\mu(i)}(p_{\mu(i)})), 0))$ for all $j \in J^+$, and the outside options are $o^+_i = o_i$ for all $i \in I^+$.[15] Since the solution (μ, p) is envy free for

[15] If $u_{i,j}(\cdot)$ is continuous then $u^{-1}_{i,j}(\cdot)$ is indeed the inverse function. More generally, it is defined for $u \in [o_i, \infty)$ by $u^{-1}_{i,j}(u) := \min_{p_j \in [r_j, \infty)}\{u_{i,j}(p_j) \leq u\}$, and is merely a one-sided inverse function satisfying $u^{-1}_{i,j}(u_{i,j}(p_j)) = p_j$.

the original problem it is also envy free for the restricted problem. It is even bidder optimal because the existence of an envy free solution (μ'', p'') for the restricted problem in which at least one bidder $i \in I^+$ has a strictly higher utility would imply the existence of an envy free solution (μ''', p''') for the original problem with this property and therefore contradict the bidder optimality of (μ, p).

Case 2.1: The solution (μ', p') is feasible for the restricted problem. That the input for the original problem is in general position implies that the input for the restricted problem is in general position. Hence Lemma 11 shows that there exists a bidder $i \in I^+$ for which $u_{i,\mu(i)}(p_{\mu(i)}) \geq u_{i,\mu'(i)}(p'_{\mu'(i)})$. This contradicts the definition of I^+.

Case 2.2: The solution (μ', p') is *not* feasible for the restricted problem. This can only happen if $p'_j < r_j^+$ for some item $j \in J^+$. Since the solution (μ', p') is feasible for the original problem this can only happen if $r_j^+ > r_j$ and, thus, $r_j^+ = \max_{i \notin I^+}(u_{i,j}^{-1}(u_{i,\mu(i)}(p_{\mu(i)})), 0)$. We cannot have $r_j^+ = 0$ as this would imply $p'_j < r_j^+ = 0$. Hence we must have $r_j^+ = u_{i,j}^{-1}(u_{i,\mu(i)}(p_{\mu(i)}))$ for some $i \in I \setminus I^+$. It follows that $p'_j < r_j^+ = u_{i,j}^{-1}(u_{i,\mu(i)}(p_{\mu(i)})) \leq u_{i,j}^{-1}(u_{i,\mu'(i)}(p'_{\mu'(i)}))$ and, thus, $u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j)$. □

Theorem 2. *If the input is in general position, then every mechanism that computes a bidder optimal solution is incentive compatible.*

Proof. For a contradiction suppose some subset of bidders $I^+ \subseteq I$ strictly benefits from misreporting their utility functions. Denote the true input by $(u_{i,j}(\cdot), r_j, o_i)$, and the falsified one by $(u'_{i,j}(\cdot)), r_j, o_i)$. Note that $u'_{i,j}(\cdot) = u_{i,j}(\cdot)$ for all $(i, j) \in I \setminus I^+ \times J$. Let (μ^*, p^*) resp. (μ', p') denote the bidder optimal solution for the true resp. falsified input. Then $I^+ = \{i \in I \mid u_{i,\mu'(i)}(p'_{\mu'(i)}) > u_{i,\mu^*(i)}(p^*_{\mu^*(i)})\}$. Note that (μ', p') is feasible for the true input $(u_{i,j}(\cdot), r_j, o_i)$ because $p'_{j_0} = 0$ and $p'_j \geq 0$ for all $j \neq j_0$.

Case 1: $I^+ = I$. Lemma 11 shows that no feasible solution (μ', p') can give all bidders a strictly higher utility than the bidder optimal solution (μ^*, p^*). This gives a contradiction.

Case 2: $I^+ \neq I$. Lemma 12 shows that if some feasible solution (μ', p') gives only some of the bidders a strictly higher utility than the bidder optimal solution (μ^*, p^*), then there must be at least one bidder $i \in I \setminus I^+$ and an item $j \in J$ for which $u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j)$. But since $i \notin I^+$ this implies that $u'_{i,\mu'(i)}(p'_{\mu'(i)}) = u_{i,\mu'(i)}(p'_{\mu'(i)}) < u_{i,j}(p'_j) = u'_{i,j}(p'_j)$ and contradicts the fact that (μ', p') is bidder optimal and therefore envy free for the falsified input $(u'_{i,j}(\cdot)), r_j, o_i)$. □

The general position concept serves as an argument to qualify inputs that are *not* in general position as degenerate. The argument is that for generic inputs, e.g., randomly generated ones, it is rather unlikely that two walks in the input graph have *exactly* the same weight. It would still be interesting, though, to be able to test whether a given input is in general position and to know how to deal with inputs that are not in general position.

The minimum is indeed contained in the set itself as we only consider right-continuous utility functions.

1.5 Conclusion and Future Work

The demand for more expressive mechanisms is reflected in the "richness of preference expression offered by businesses as diverse as matchmaking sites, sites like Amazon and NetFlix, and services like Google's AdSense" [7]. Standard mechanisms often do not meet this demand. Providing mechanisms that do meet this demand and that at the same time (1) guarantee the existence of a stable solution, (2) are computationally tractable, and (3) have good incentive properties is one of the major challenges that the field of algorithmic mechanism design is currently facing.

In this chapter we contributed to this agenda by considering the domain of *multi-item auctions* with *unit demand* and by providing the most expressive mechanism for this setting so far. This mechanism, which can be seen as a generalization of the General Auction Mechanism of [1], can handle *piece-wise linear* utility functions with *non-identical slopes* and *multiple discontinuities*. These utility functions allow the bidders to explicitly specify *conversion rates* (enabling, e.g., per-click auctions that are simultaneously envy free for bidders with per-click *and* per-impression valuations) and a variety of *soft* and *hard budget constraints* (which, e.g., arise when bidders have a limited amount of cash and have to take out loans).

An interesting direction for future work would be to devise even more expressive mechanisms for the domain studied here, or to devise expressive mechanisms for more general domains (e.g., one-to-many or many-to-many matchings).

Bibliography

1. G. Aggarwal, S. Muthukrishnan, D. Pál, and M. Pál. General auction mechanism for search advertising. In *Proceedings of the 18th International World Wide Web Conference*, pages 241–250, 2009.

2. S. Alaei, K. Jain, and A. Malekian. Competitive equilibrium in two sided matching markets with general utility functions. *SIGecom Exchanges*, 10(2):Article 8, 2011.

3. A. Alkan. Existence and computation of matching equilibria. *European Journal of Political Economy*, 5(2-3):285–296, 1989.

4. M. Alkan. Equilibrium in a matching market with general preferences. In M. Majumdar, editor, *Equilibrium and Dynamics. Essays in Honour of David Gale*, pages 1–16. The MacMillan Press Ltd, London, 1992.

5. I. Ashlagi, M. Braverman, and A. Hassidim. Ascending unit demand auctions with budget limits. *Working paper*, 2009.

6. I. Ashlagi, M. Braverman, A. Hassidim, R. Lavi, and M. Tennenholtz. Position auctions with budgets: Existence and uniqueness. *The B.E. Journal of Theoretical Economics (Advances)*, 10(1):Article 20, 2010.

7. M. Benisch, N. Sadeh, and T. Sandholm. A theory of expressiveness in mechanisms. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, pages 17–23, 2008.

8. N. Chen, X. Deng, and A. Ghosh. Competitive equilibria in matching markets with budgets. *SIGecom Exchanges*, 9(1):Article 5, 2010.

9. E. H. Clarke. Multipart pricing of public goods. *Public Choice Journal*, 11(1):17–33, 1971.

10. G. Demange and D. Gale. The strategy structure of two-sided matching markets. *Econometrica*, 53(4):873–888, 1985.

11. G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–72, 1986.

12. X. Deng and Y. Du. The computation of approximate competitive equilibrium is PPAD-hard. *Information Processing Letters*, 108(6):369–373, 2008.

13. P. Dütting, M. Henzinger, and I. Weber. Sponsored search, market equilibria, and the Hungarian Method. *Information Processing Letters*, 113(3):67–73, 2013.

14. P. Dütting, M. Henzinger, and I. Weber. Bidder optimal assignments for general utilities. *Theoretical Computer Science*, 478:22–32, 2013.

15. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

16. J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM*, 19(2):248–264, 1972.

17. T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

18. P. Hall. On representatives of subsets. *London Mathematical Society*, 10:26–30, 1935.

19. D. Kempe, A. Mu'alem, and M. Salek. Envy-free allocations for budgeted bidders. In *Proceedings of the 5th International Workshop on Internet and Network Economics*, volume 5929 of *Lecture Notes in Computer Science*, pages 537–544, 2009.

20. H. W. Kuhn. The Hungarian Method for the assignment problem. *Naval Research Logistics*, 2:83–97, 1955.

21. H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479, 1983.

22. A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, New York, 1998.

23. L. S. Shapley and M. Shubik. The assignment game: The core I. *International Journal of Game Theory*, 1(1):111–130, 1971.

24. N. Tomizawa. On some techniques useful for solution of transportation network problems. *Networks*, 1(2):173–194, 1971.

25. H. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6): 1163–1178, 2007.

26. V. Vazirani and M. Yannakakis. Market equilibrium under separable, piecewise-linear, concave utilities. In *Proceedings of the 1st Symposium on Innovations in Computer Science*, pages 156–165, 2010.

27. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Finance*, 16(1): 8–27, 1961.

2 Simplicity-Expressiveness Tradeoffs in Mechanism Design

2.1 Introduction

The revelation principle states that any equilibrium outcome of any mechanism can be obtained in a truthful equilibrium of a direct-revelation mechanism, i.e., an equilibrium in which agents truthfully reveal their types. In practice, however, the assumptions underlying the revelation principle often fail: other equilibria may exist besides the truthful one, or computational limitations may interfere. As a consequence, many practical mechanisms use a restricted message space. This motivates the study of the tradeoffs involved in choosing simplified mechanisms, which can sometimes bring benefits in precluding bad or promoting good equilibria, and other times impose costs on welfare and revenue. Despite their practical importance, these tradeoffs are currently only poorly understood.

In sponsored search auctions, and adopting a complete information analysis, allowing every agent i to submit a valuation v_{ij} for each slot j means that both the Vickrey-Clarke-Groves (VCG) mechanism and the Generalized Second Price (GSP) mechanism always admit an efficient Nash equilibrium with zero revenue [23]. In fact, if agents face a small cost for submitting a positive bid, then all Nash equilibria yield zero revenue. If instead agent i is asked for a single bid b_i, and his bid for slot j is derived by multiplying b_i with a slot-specific click-through rate α_j, the zero-revenue equilibria are eliminated. More surprisingly, this simplification does not introduce new equilibria (even if α_j is not correct for every agent i), so minimum revenue over all Nash equilibria is strictly greater for the simplification than for the original mechanism. Moreover, if valuations can actually be decomposed into an agent-specific valuation v_i and click-through rates α_j, the simplification still has an efficient Nash equilibrium. Milgrom [23] concluded that simplification can be beneficial and need not come at a cost.

For combinatorial auctions, and adopting an incomplete information analysis, the maximum social welfare over all outcomes of a mechanism strictly increases with expressiveness, for a particular measure of expressiveness based on notions from computational learning theory [4]. Implicit in the result of Benisch et al. [4] is the conclusion that more expressiveness is generally desirable, as it allows a mechanism to achieve a more efficient outcome in more instances of the problem.

Each of these results tells only part of the story. While Milgrom's results on the benefits of simplicity are developed within an equilibrium framework and can be extended beyond sponsored search auctions, they critically require settings with complete information amongst agents in order to preclude bad equilibria while retaining good ones. In particular, Milgrom does not consider the potential loss in efficiency or revenue that can occur when agents are ignorant of each other's valuations when deciding how to bid within a simplified bidding language. Benisch et al., on the other hand, develop their results on the benefits of expressiveness (and thus the cost of simplicity) in an incomplete information context, but largely in the absence of equilibrium considerations.[1] In particular, these authors do not consider the potential problems that can occur due to the existence of bad equilibria in expressive mechanisms.

2.1.1 Our Contribution

Our contribution is twofold. On a conceptual level, we analyze how different properties of a simplification affect the set of equilibria of a mechanism in both complete and incomplete information settings and argue that well-chosen simplifications can have a positive impact on the set of equilibria as a whole; either by precluding undesirable equilibria or by promoting desirable equilibria. We thus extend Milgrom's emphasis on simplification as a tool that enables equilibrium selection. On a technical level, we analyze simplified mechanisms for sponsored search auctions with complete information and combinatorial auctions with incomplete information.

An important property when analyzing the impact of simplification on the set of equilibria is *tightness* [23], which requires that no additional equilibria are introduced. We observe that tightness can be achieved equally well in complete and incomplete information settings and give a sufficient condition.[2] Complementary to tightness is a property we call *totality*, which requires all equilibria of the original mechanism to be preserved. To the end of equilibrium selection, totality needs to be relaxed. Particular relaxations we consider are the preservation of the VCG outcome, i.e., the outcome obtained in the dominant strategy equilibrium of the fully expressive VCG mechanism, and the existence of an equilibrium with a certain amount of social welfare or revenue relative to the VCG outcome. In addition, one might require that the latter property holds for *every* equilibrium of the simplified mechanism.

In the context of sponsored search, one reason to prefer a simplification is to preclude the zero-revenue equilibria discussed above. Another interesting property of a simplified GSP

[1] The equilibrium analysis that Benisch et al. provide is in regard to identifying a particular mechanism design in which the maximum social welfare achievable in any outcome can be achieved in a particular Bayes-Nash equilibrium.

[2] This condition was already considered by Milgrom, but only in the complete information case.

mechanism is that it preserves the VCG outcome even when the assumed click-through rates α_j are inexact. Recognizing that this claim cannot be made for the VCG mechanism under the same simplification, Milgrom [23] uses this as an argument for the superiority of the GSP mechanism. But, this result that GSP is Vickrey-preserving requires an unnatural condition on the relationship between the assumed click-through rates and prices and thus agents' bids, and moreover does not preclude alternate simplifications of VCG that succeed in being Vickrey-preserving. In addition, we observe that the simplifications can still suffer arbitrarily low revenue in some equilibrium, in comparison with the VCG revenue. In our analysis of sponsored search, we identify a simplified GSP mechanism that preserves the VCG outcome without requiring any knowledge of the actual click-through rates, precludes zero-revenue equilibria, and always recovers at least half of the VCG payments for all slots but the first. For simplified VCG mechanisms, we obtain a strong negative result: *every* simplification of VCG that supports the (efficient) VCG outcome in some equilibrium also has an equilibrium in which revenue is arbitrarily smaller than in the VCG outcome.

In the context of combinatorial auctions, paradigmatic of course of settings with incomplete information, we first recall the previous observations by Holzman et al. [16, 17] in regard to the existence of multiple, non-truthful, *ex-post* equilibria of the VCG mechanism, each of which offers different welfare and revenue properties. In particular, if one assumes that participants will select equilibria with a particular (maximum) number of bids, social welfare can differ greatly among the different equilibria, revenue can be zero for some of them, and the existence of multiple Pareto optimal equilibria can make equilibrium selection hard to impossible for participants. Focusing again on *tight* simplifications, we connect the analysis of Holzman et al. [16, 17] with tightness, by establishing that a simplification is tight if and only if bids are restricted to a subset Σ of the bundles with a *quasi-field* structure [16, 17], with values for the other bundles derived as the maximum value of any contained bundle. Through insisting on a tight simplification, we ensure that the worst-case behavior is no worse than that of the fully expressive VCG mechanism, even when Σ (although simplified) contains too many bundles for agents to bid on all of them. Moreover, using a quasi-field simplification ensures that agents do not experience regret with respect to the bidding language, in the sense of wanting to send a message *ex post* that was precluded. Finally, as any restriction of the bids to a subset of the bundles, restricting the bids to a quasi-field Σ yields a mechanism that is maximal in range [24] and makes it a dominant strategy equilibrium for the agents to bid truthfully on these bundles. Simplification thus enables the mechanism designer to guide equilibrium selection, and our results suggest that the presence or absence of such guidance can have a significant impact on the economic properties of the mechanism.

The informational assumptions underlying our analysis are crucial, and the amount of information available to the agents plays an important role for the tradeoff between simplicity and expressiveness. In the sponsored search setting, both the existence of a zero-revenue equilibrium in the expressive mechanism, and the existence of the desirable equilibrium in the simplification, rely on the assumption of complete information. In other words, agents can on one hand use information about each others' types to coordinate and harm the auctioneer, but on the other hand, the same information guarantees that simplified mechanisms retain the desirable equilibria of the expressive mechanism. In combinatorial auctions the contrast is equally stark: while bids on every single package may be required to sustain an efficient equilibrium in the incomplete information setting, we show that such an equilibrium can be obtained with a number of bids that is quadratic in the number of agents, and potentially exponentially smaller than the number of bundles, given that agents have complete information.

2.1.2 Related Work

Several authors have criticized the revelation principle because it does not take computational aspects of mechanisms into account. In this context, Conitzer and Sandholm [10] consider sequential mechanisms that reduce the amount of communication, and non-truthful mechanisms that shift the computational burden of executing the mechanism, and the potential loss when it is executed suboptimally, from the designer to the agents. Hyafil and Boutilier [18, 19] propose to circumvent computational problems associated with direct type revelation via the automated design of partial-revelation mechanisms, and in particular study approximately incentive compatible mechanisms that do not make any assumptions about agents' preferences. This approach is very general, but also hard to analyze theoretically, with complex, regret-based algorithms.

Blumrosen et al. [6] and Feldman and Blumrosen [12] consider settings with one-dimensional types and ask how much welfare and revenue can be achieved by mechanisms with a bounded message space. By contrast, we study mechanisms with message spaces that grow in some parameter of the problem and may even have infinite size, and obtain results both for one-dimensional and multi-dimensional types.

A different notion of simplicity of a mechanism was considered by Babaioff and Roughgarden [3]: the authors show that among all payment rules that guarantee an efficient equilibrium when ranking agents according to their bids, the GSP payment rule is optimally simple in the sense that prices depend on bids in a minimal way.

Shakkottai et al. [26] study the tradeoff between simplicity and revenue in the context of pricing rules for communication networks and define the "price of simplicity" as the ratio between the revenue of a very simple pricing rule and the maximum revenue that can be obtained.

2.2 Preliminaries

A mechanism design problem is given by a set $N = \{1, 2, \dots, n\}$ of *agents* that interact to select an element from a set Ω of *outcomes*. Agent $i \in N$ is associated with a *type* θ_i from a set Θ_i of possible types, representing private information held by this agent. We write $\theta = (\theta_1, \dots, \theta_n)$

for a profile of types for the different agents, $\Theta = \prod_{i \in N} \Theta_i$ for the set of possible type profiles, and $\theta_{-i} \in \Theta_{-i}$ for a profile of types for all agents but i. Each agent $i \in N$ further employs preferences over Ω, represented by a *valuation function* $v_i : \Omega \times \Theta_i \to \mathbb{R}$. The quality of an outcome $o \in \Omega$ is typically measured in terms of its social welfare, which is defined as the sum $\sum_{i \in N} v_i(o, \theta_i)$ of agents' valuations. An outcome that maximizes social welfare is also called *efficient*.

A *mechanism* is given by a tuple (N, X, f, p), where $X = \prod_{i \in N} X_i$ is a set of *message profiles*, $f : X \to \Omega$ is an *outcome function*, and $p : X \to \mathbb{R}^n$ is a *payment function*. In this chapter, we mostly restrict our attention to direct mechanisms, i.e., mechanisms where $X_i \subseteq \Theta_i$ for every $i \in N$. A direct mechanism (N, X, f, p) with $X = \Theta$ is called *efficient* if for every $\theta \in \Theta$, $f(\theta)$ is an efficient outcome. Just as for type profiles, we write $x_{-i} \in X_{-i}$ for a profile of messages by all agents but i. We assume quasilinear preferences, i.e., the utility of agent i given a message profile $x \in X$ is $u_i(x, \theta_i) = v_i(f(x), \theta_i) - p_i(x)$. The *revenue* achieved by mechanism (N, X, f, p) for a message profile $x \in X$ is $\sum_{i \in N} p_i(x)$.

Game-theoretic reasoning is used to analyze how agents interact with a mechanism, a desirable criterion being stability according to some game-theoretic solution concept. We consider two different settings. In the *complete information* setting, agents are assumed to know the type of every other agent. A strategy of agent i in this setting is a function $s_i : \Theta \to X_i$. In the *(strict) incomplete information* setting, agents have no information, not even distributional, about the types of the other agents. A strategy of agent i in this setting thus becomes a function $s_i : \Theta_i \to X_i$.

The two most common solution concepts in the complete information setting are dominant strategy equilibrium and Nash equilibrium. A *dominant strategy equilibrium* consists of a dominant strategy $s_i : \Theta \to X_i$ for each agent $i \in N$. A strategy $s_i : \Theta \to X_i$ is a *dominant strategy* for agent $i \in N$ if for every $\theta \in \Theta$, every $x_{-i} \in X_{-i}$, and every $x_i \in X_i$,

$$u_i((s_i(\theta), x_{-i}), \theta_i) \geq u_i((x_i, x_{-i}), \theta_i).$$

A profile $s \in \prod_{i \in N} s_i$ of strategies $s_i : \Theta \to X_i$ is a *Nash equilibrium* if for every $\theta \in \Theta$, every $i \in N$, and every $s'_i : \Theta \to X_i$,

$$u_i((s_i(\theta), s_{-i}(\theta)), \theta_i) \geq u_i((s'_i(\theta), s_{-i}(\theta)), \theta_i).$$

The existence of a dominant strategy $s_i : \Theta \to X_i$ always implies the existence of a dominant strategy $s'_i : \Theta_i \to X_i$ that does not depend on the types of other agents. The solution concept of dominant strategy equilibrium thus carries over directly to the incomplete information setting. Formally, a *dominant strategy equilibrium* in this setting consists of a dominant strategy $s_i : \Theta_i \to X_i$ for each agent $i \in N$. A strategy $s_i : \Theta_i \to X_i$ is a *dominant strategy* for agent $i \in N$ if for every $\theta_i \in \Theta_i$, every $x_{-i} \in X_{-i}$, and every $x_i \in X_i$,

$$u_i((s_i(\theta_i), x_{-i}), \theta_i) \geq u_i((x_i, x_{-i}), \theta_i).$$

The appropriate variant of the Nash equilibrium concept in the incomplete information setting is that of an ex-post equilibrium. A profile $s \in \prod_{i \in N} s_i$ of strategies $s_i : \Theta_i \to X_i$ is an *ex-post equilibrium* if for every $\theta \in \Theta$, every $i \in N$, and every $s'_i : \Theta_i \to X_i$,

$$u_i((s_i(\theta_i), s_{-i}(\theta_{-i})), \theta_i) \geq u_i((s'_i(\theta_i), s_{-i}(\theta_{-i})), \theta_i).$$

We conclude this section with a direct mechanism due to Vickrey [28], Clarke [9], and Groves [15]. This mechanism starts from an efficient outcome function f and computes each agent's payment according to the total value of the other agents, thus aligning his interests with that of society. Formally, mechanism (N, X, f, p) is called Vickrey-Clarke-Groves (VCG) mechanism[3] if $X = \Theta$, f is efficient, and

$$p_i(\theta) = \max_{o \in \Omega} \sum_{j \neq i} v_j(o, \theta_j) - \sum_{j \neq i} v_j(f(\theta), \theta_j).$$

In the VCG mechanism, revealing types $\theta \in \Theta$ truthfully is a dominant strategy equilibrium [15]. We will refer to the resulting outcome as the VCG outcome for θ, and write $R^{VCG}(\theta)$ for the revenue obtained in this outcome.

2.3 Simplifications

Our main object of study in this chapter are simplifications of a mechanism obtained by restricting its message space. Consider a mechanism $M = (N, X, f, p)$. A mechanism $\hat{M} = (N, \hat{X}, \hat{f}, \hat{p})$ will be called a *simplification* of M if $\hat{X} \subseteq X$, $\hat{f}|_{\hat{X}} = f|_{\hat{X}}$, and $\hat{p}|_{\hat{X}} = p|_{\hat{X}}$. We will naturally be interested in the set of outcomes that can be obtained in equilibrium, both in the original mechanism M and the simplified mechanism $\hat{M}$.[4]

Milgrom [23] defines a property he calls tightness, which requires that the simplification does not introduce any additional equilibria. More formally, simplification $\hat{M}$ will be called *tight* if every equilibrium of $\hat{M}$ is an equilibrium of M.[5] Tightness ensures that the simplified mechanism is at least as good as the original one with respect to the worst outcome obtained in any equilibrium. It does not by itself protect good equilibrium outcomes, and we will in fact see examples of tight simplifications that eliminate *all* ex-post equilibria. A property that will be useful in the following is a variant of Milgrom's outcome closure for exact equilibria. It requires that for every agent, and for every choice of messages from the restricted message spaces of the other agents, it is optimal for the agent to choose a message from his restricted message space. More formally, a simplification $(N, \hat{X}, \hat{f}, \hat{p})$ of a mechanism (N, X, f, p) satisfies

[3] Actually, we consider a specific member of a whole family of VCG mechanisms, namely the one that uses the Clarke pivot rule.

[4] In the following, we will simply talk about equilibria without making a distinction between the different equilibrium notions. Unless explicitly noted otherwise, our results concern Nash equilibria in the complete information case and ex-post equilibria in the incomplete information case.

[5] Milgrom [23] considers a slightly stronger notion of tightness defined with respect to (pure-strategy) ϵ-Nash equilibria.

outcome closure if for every $\theta \in \Theta$, every $i \in N$, every $\hat{x}_{-i} \in \hat{X}_{-i}$, and every $x_i \in X_i$ there exists $\hat{x}_i \in \hat{X}_i$ such that $u_i((\hat{x}_i, \hat{x}_{-i}), \theta_i) \geq u_i((x_i, \hat{x}_{-i}), \theta_i)$. This turns out to be sufficient for tightness in both the complete and incomplete information case.[6]

Proposition 1 (Milgrom [23])**.** *Every simplification that satisfies outcome closure is tight with respect to both Nash and ex-post equilibria.*

Proof. Fix $\theta \in \Theta$. Consider a mechanism $M = (N, X, f, p)$, a simplification $\hat{M} = (N, \hat{X}, \hat{f}, \hat{p})$ that satisfies outcome closure, and an equilibrium $\hat{x}$ of $\hat{M}$. Assume for contradiction that $\hat{x}$ is not an equilibrium of M. Then, for some $i \in N$, there exists $x'_i \in X_i$ such that $u_i((x'_i, \hat{x}_{-i}), \theta_i) > u_i(\hat{x}, \theta_i)$. Since $\hat{M}$ satisfies outcome closure, there further exists $\hat{x}'_i \in \hat{X}_i$ such that $u_i((\hat{x}'_i, \hat{x}_{-i}), \theta_i) \geq u_i((x'_i, \hat{x}_{-i}), \theta_i)$. It follows that $u_i((\hat{x}'_i, \hat{x}_{-i}), \theta_i) > u_i(\hat{x}, \theta_i)$, which contradicts the assumption that $\hat{x}$ is an equilibrium of $\hat{M}$. □

One way to guarantee good behavior in the *best case* is by requiring that a simplification $\hat{M}$ preserves all equilibria of the original mechanism M, in the sense that for every Nash equilibrium of M, there exists a Nash equilibrium of $\hat{M}$ that yields the same outcome and payments. We will call a simplification satisfying this property *total*. A useful property in this context is outcome reducibility, which requires that there exists a way of mapping messages of the original mechanism to messages of the simplification such that (i) outcomes and payoffs are preserved, and such that (ii) for every agent the utility that he can achieve with a message from the original message space against messages of the other agents from the original message space is at least as high as the utility that he can achieve with a message from the restricted message space against the corresponding mapped messages of the other agents. More formally, a simplification $(N, \hat{X}, \hat{f}, \hat{p})$ of a mechanism (N, X, f, p) satisfies *outcome reducibility* if there exists a mapping $h : X \to \hat{X}$ such that (i) for every $x \in X$, $f(x) = f(h(x))$ and $p(x) = p(h(x))$ and (ii) for every $i \in N$, every $x \in X$, and every $\hat{x}'_i \in \hat{X}_i$, there exists $x'_i \in X_i$ such that $u_i((x'_i, x_{-i}), \theta_i) \geq u_i((\hat{x}'_i, h_{-i}(x)), \theta_i)$. Outcome reducibility is sufficient for totality, but only with respect to Nash equilibria.

Proposition 2. *Every simplification that satisfies outcome reducibility is total with respect to Nash equilibria.*

Proof. Fix $\theta \in \Theta$. Consider a mechanism $M = (N, X, f, p)$, a simplification $\hat{M} = (N, \hat{X}, \hat{f}, \hat{p})$ that satisfies outcome reducibility, and a Nash equilibrium x of M. Assume for contradiction that $h(x) \in \hat{X}$ is not a Nash equilibrium of $\hat{M}$. Then, for some $i \in N$, there exists $\hat{x}'_i \in \hat{X}_i$ such that $u_i((\hat{x}'_i, h_{-i}(x)), \theta_i) > u_i(h(x), \theta_i)$. Since $\hat{M}$ satisfies outcome reducibility, there further exists $x'_i \in X_i$ such that $u_i((x'_i, x_{-i}), \theta_i) \geq u_i((\hat{x}'_i, h_{-i}(x)), \theta_i)$. It follows that $u_i((x'_i, x_{-i}), \theta_i) > u_i(h(x), \theta_i) = u_i(x, \theta_i)$, which contradicts the assumption that x is a Nash equilibrium of M. □

[6] Milgrom stated the result only for complete information, but the proof goes through for incomplete information as well.

To the end of equilibrium selection totality clearly needs to be relaxed, by requiring that only certain desirable outcomes are preserved. A typical desirable outcome in many settings is the VCG outcome. This outcome has some shortcomings in fully general combinatorial auction domains [2], but it remains of significant interest in settings with unit-demand preferences, such as sponsored search. We will call simplification $\hat{M}$ *Vickrey-preserving* if for every $\theta \in \Theta$, it has an equilibrium that yields the VCG outcome for θ.

2.4 Sponsored Search Auctions

In sponsored search (see, e.g., [21]), the agents compete for elements of a set $S = \{1,\dots,k\}$ of slots, where $k \leq n$. Each outcome corresponds to a one-to-one assignment of agents to slots, i.e., $\Omega \subseteq \{1,\dots,n\}^N$ such that $o_i \neq o_j$ for all $o \in \Omega$ and $i, j \in \{1,\dots,n\}$ with $i \neq j$. We will assume that $v(o,\theta_i) = 0$ if $o_i > k$, and that there are no externalities, i.e., $v_i(o,\theta_i) = v_i(o',\theta_i)$ if $o_i = o'_i$. In slight abuse of notation, we will write $v_i(j,\theta_i)$ for the valuation of agent i for slot j.

We consider simplifications of two mechanisms, the Vickrey-Clarke-Groves (VCG) mechanism and the Generalized Second Price (GSP) mechanism, and analyze their behavior for different spaces of type profiles, which we denote by Θ, $\Theta^>$, and Θ^α. In Θ, valuations can be arbitrary non-negative numbers. $\Theta^>$ adds the restriction that valuations are strictly decreasing, i.e., $v_i(j,\theta_i) > v_i(j+1,\theta_i)$ for every $\theta \in \Theta^>$, $i \in N$, and $j \in \{1,\dots,k-1\}$. Valuations in Θ^α are assumed to arise from "clicks" associated with each slot and a valuation per click. In other words, there exists a fixed click-through rate vector $\alpha \in \mathbb{R}^n_> = \{\alpha' \in \mathbb{R}^n_{\geq 0} : \alpha'_i > \alpha'_j \text{ if } i < j \text{ and } i \leq k\}$, which may or may not be known to the mechanism, and $v_i(j,\theta_i) = \alpha_j \cdot v_i(\theta_i)$ for some $v_i(\theta_i) \in \mathbb{R}_{\geq 0}$. Thus, $\alpha_j = 0$ if $j > k$, and it will be convenient to assume that $\alpha_1 = 1$. We finally define $\Psi = \bigcup_{\alpha \in \mathbb{R}^n_>} \Theta^\alpha$, and observe that $\Psi \subset \Theta^> \subset \Theta$.

The message an agent $i \in N$ submits to the mechanism thus corresponds to a vector of bids $x_i \in \mathbb{R}^n$ specifying a bid $x_{i,j} \in \mathbb{R}$ for each slot $j \in S$. Given a message profile $x \in X$, the VCG mechanism assigns each agent i a slot $f_i(x) = o_i$ so as to maximize $\sum_i x_{i,o_i}$, and charges that agent $p_i(x) = \max_{o' \in \Omega} \sum_{j \neq i} x_{j,o'_j} - \sum_{j \neq i} x_{j,o_j}$. The GSP mechanism is defined via a sequence of second-price auctions for slots 1 through k: slot j is assigned to an agent i with a maximum bid for that slot at a price equal to the second highest bid, both with respect to the set of agents who have not yet been assigned a slot, i.e., $f_i(x) = o_i$ such that $x_{i,j} = \max_{i' \in N: o_{i'} \geq j} x_{i',j}$ and $p_i(x) = \max_{i' \in N: o_{i'} > j} x_{i,j}$.

2.4.1 Envy Freeness and Efficiency

The original analysis of GSP due to Edelman et al. [11] and Varian [27] focuses on equilibria that are "locally envy free."[7] Assume that $\theta \in \Theta^\alpha$, and consider an outcome in which agent i is assigned slot i, for all $i \in N$. Such an outcome is called locally envy free if, in addition to being a Nash equilibrium, no agent could increase his utility by exchanging bids with

[7] Varian calls these equilibria "symmetric equilibria."

the agent assigned the slot directly above him, i.e., if for every $i \in \{2,\ldots,n\}$, $\alpha_i \cdot v_i(\theta_i) - p_i \geq \alpha_{i-1} \cdot v_i(\theta_i) - p_{i-1}$. Restricting attention to envy free equilibria immediately solves all revenue problems: as Edelman et al. point out, revenue in any locally envy free equilibrium is at least as high as that in the dominant-strategy equilibrium of the VCG auction. Conversely, Milgrom's observation concerning zero-revenue equilibria contains an implicit critique of the assumptions underlying the restriction to equilibria that are envy free. It will be instructive to make this critique explicit.

Edelman et al. argue that an equilibrium where some agent i envies some other agent j assigned the next higher slot is not a reasonable rest point of the bidding process because agent i might increase the price paid by agent j without the danger of harming his own utility should j retaliate. There are two problems with this line of reasoning. First, it is not clear why agent j should retaliate, especially if he is worse off by doing so. Second, agent i might in fact have a very good reason *not* to increase the price paid by j, like a desire to keep prices low in the long run through tacit collusion.

We may instead ask under what conditions there exists a bid $x_{i,j}$ for agent i on slot $j < i$ such that (i) agent j is forced out of slot j, in the sense that it becomes a better response for j to underbid i, and (ii) agent i is strictly better off after this response by j than at present. This is the case exactly when

$$\alpha_i \cdot v_j(\theta_j) - p_i > \alpha_j \cdot v_j(\theta_j) - x_{i,j} \qquad \text{and} \qquad \alpha_j \cdot v_i(\theta_i) - x_{i,j} > \alpha_i \cdot v_i(\theta_i) - p_i, \tag{2.1}$$

where p_i is the price currently paid by agent i. The second inequality assumes that j will respond by bidding just below the bid $x_{i,j}$ of agent i, such that this becomes i's new price.

An equilibrium that does not allow deviations as above will be called *two-step stable*. It turns out that two-step stability exactly characterizes the set of efficient equilibria.

Proposition 3. *Let $\alpha \in \mathbb{R}^n_>$, $\theta \in \Theta^\alpha$. Then, an equilibrium is two-step stable if and only it is efficient.*

Proof. Consider agents $i > j$. Suppose agent i is assigned slot i and agent j is assigned slot j. Rewriting (2.1), it must hold that

$$x_{i,j} > (\alpha_j - \alpha_i) \cdot v_j(\theta_j) + p_i \qquad \text{and} \qquad x_{i,j} < (\alpha_j - \alpha_i) \cdot v_i(\theta_i) + p_i.$$

Clearly, a bid $x_{i,j}$ with this property exists if and only if $v_i(\theta_i) > v_j(\theta_j)$. In turn, agents i and j such that $v_i(\theta_i) > v_j(\theta_j)$ exist if and only if the current assignment is inefficient. □

This result provides a very strong argument against inefficient equilibria as rest points of the bidding process, much stronger than the argument against equilibria that are not envy free. In the context of sponsored search auctions we will therefore restrict our attention to efficient equilibria. It is worth noting at this point that the set of efficient equilibria forms a strict superset of the set of locally envy free equilibria, and in particular contains the zero-revenue equilibria identified by Milgrom and discussed next.

2.4.2 Comments on Milgrom's Analysis

Milgrom [23] observed that for every profile of agent types, both expressive VCG and expressive GSP have an efficient equilibrium that yields zero revenue, and that all equilibria yield zero revenue if there is a small cost associated with submitting a positive bid. To alleviate this fact, he proposed to restrict the message space of both VCG and GSP to $\hat{X} = \{(\alpha_1 \cdot b_i, \ldots, \alpha_k \cdot b_i) : b_i \in \mathbb{R}_{\geq 0}\}$ for some $\alpha \in \mathbb{R}^n_>$.

The following proposition summarizes our knowledge about the resulting simplifications, which we will refer to as α-VCG and α-GSP. Most of these observations were already proved or at least claimed by Milgrom, but a proof of the proposition is given for the sake of completeness.

Proposition 4. *Let $\alpha \in \mathbb{R}^n_>$. Then, α-GSP and α-VCG are tight on Θ, have positive revenue on $\Theta^>$ if $n, k \geq 2$, and are Vickrey-preserving on Θ^α.*

Proof. First we prove that α-GSP and α-VCG are tight on Θ. By Proposition 1 it suffices to show that for every $\theta \in \Theta$, every $i \in N$, every $\hat{x}_{-i} \in \hat{X}_{-i}$ and every $x_i \in X_i$ there exists $\hat{x}_i \in \hat{X}_i$ such that $u_i((\hat{x}_i, \hat{x}_{-i}), \theta_i) \geq u_i((x_i, \hat{x}_{-i}), \theta)$. Denote the outcome of GSP resp. VCG on $(x_i, \hat{x}_{-i})$ by o. For $l \neq i$ let $b_l \in \mathbb{R}_{\geq 0}$ be such that $\hat{x}_l = (\alpha_1 \cdot b_l, \ldots, \alpha_k \cdot b_l)$. Let b_i be the o_i-th highest value among the b_l's. Then in the outcome o' of α-GSP and α-VCG on $(\hat{x}_i, \hat{x}_{-i})$ we have $o'_i = o_i$ and, thus, $v_i(o'_i, \theta_i) = v_i(o_i, \theta_i)$. Since i's payment $p_i(o_i)$ and $p_i(o'_i)$ in GSP resp. VCG and α-GSP resp. α-VCG only depends on $\hat{x}_{-i}$ and is therefore the same, we conclude that $u_i((\hat{x}_i, \hat{x}_{-i}), \theta_i) \geq u_i((x_i, \hat{x}_{-i}), \theta)$.

Next we show that α-GSP and α-VCG have positive revenue on $\Theta^>$ for $n, k \geq 2$. Suppose $\hat{x}$ is an equilibrium of α-GSP resp. α-VCG. For $i \in N$ let $b_i \in \mathbb{R}_{\geq 0}$ be such that $\hat{x}_i = (\alpha_1 \cdot b_i, \ldots, \alpha_k \cdot b_i)$. Renumber the agents by non-increasing b_i. The revenue achieved by α-GSP resp. α-VCG is $\sum_i \alpha_i \cdot b_{i+1}$ resp. $\sum_i \sum_{j>i} (\alpha_{j-1} - \alpha_j) \cdot b_j$, where the sums are over all $i, j \leq \min(n, k)$ with $b_i, b_j > 0$. Hence in both, α-GSP and α-VCG, we can only have zero revenue if $b_i = 0$ for all $i > 1$. In this case all agents but the first remain unassigned and any agent $i > 1$ can bid $0 < b'_i < b_1$ to be assigned slot 2 at price 0. Since $\theta_i \in \Theta^>_i$ we have $v_i(2, \theta_i) > 0$ and, thus, agent i's utility would strictly increase. We conclude that in both, α-GSP and α-VCG, the revenue associated with $\hat{x}$ must be strictly positive.

Finally, we show that α-GSP and α-VCG are Vickrey compatible on Θ^α. For this denote the outcome and payments computed by VCG for the true types $\theta \in \Theta^\alpha$ by o and p. Recall that $i < j$ implies that $p_i > p_j$. We have to argue that there are equilibria $\hat{x} \in \hat{X}$ of α-GSP and α-VCG in which the outcome and payments are identical to o and p.

For α-GSP we construct $\hat{x} \in \hat{X}$ as follows: 1. Renumber the agents by the slot they are assigned to in o. 2. Set $b_1 = v_1(1,\theta_1)$ and for $i > 1$ set $b_i = \alpha_{i-1}^{-1} \cdot p(i-1)$. 3. For all i let $\hat{x}_i = (\alpha_1 \cdot b_i, \ldots, \alpha_k \cdot b_i)$. We have chosen $\hat{x} \in \hat{X}$ such that the outcome and payments computed by α-GSP on $\hat{x}$ are identical to o and p. To see that $\hat{x}$ is an equilibrium of α-GSP observe that if agent i deviates from $\hat{x}$ to win slot j, then the price $p''(j)$ that he would have to pay for slot j is at least as large as the price $p'(j) = p(j)$ of slot j in the outcome computed by α-GSP resp. VCG on $\hat{x}$ resp. θ. That is, if i would strictly benefit from the deviation that gives him slot j, then we would have $v_i(j,\theta_i) - p(j) \geq v_i(j,\theta_i) - p''(j) > v_i(i,\theta_i) - p'(i) = v_i(i,\theta_i) - p(i)$. But this would contradict the fact that the outcome and payments computed by VCG are envy free [22].

For α-VCG we can use $\hat{x} = \theta \in \hat{X}$. Applying α-VCG to $\hat{x}$ gives, of course, the same assignment and payments as applying VCG to θ. To see that $\hat{x}$ is an equilibrium of α-VCG observe that any beneficial deviation from $\hat{x}$ in α-VCG would also be a beneficial deviation from θ in VCG and, thus, would contradict the fact that truthtelling is a dominant strategy of VCG. □

An additional observation that we make, in regard to the ability of the simplifications α-VCG and α-GSP to eliminate zero-revenue equilibria, is that there exist type profiles for which the minimum equilibrium revenue can be *arbitrarily small* compared to the revenue obtained in the VCG outcome.

Theorem 1. *Let $\epsilon, r > 0$. Then there exist $\alpha \in \mathbb{R}^n_>$ and $\theta \in \Theta^\alpha$ such that $R^{VCG}(\theta) \geq r$ and α-VCG has an equilibrium with revenue at most ϵ. Similarly, there exist $\alpha \in \mathbb{R}^n_>$ and $\theta \in \Theta^\alpha$ such that $R^{VCG}(\theta) \geq r$ and α-GSP has an equilibrium with revenue at most ϵ.*

Proof. We consider a setting with three agents and three slots. The construction can easily be extended to an arbitrary number of agents and slots.

For α-VCG, let $v_1(\theta_1) = v_2(\theta_2) = r+1$ and let $v_3(\theta_3) = \epsilon$. Let $\alpha_1 = 1$, $\alpha_2 = 1/(r+1)$, and $\alpha_3 = 1/(2r+2)$. It is easily verified that the bids $b_1 = r+1$ and $b_2 = b_3 = \epsilon$ form an equilibrium of α-VCG. Given these bids, α-VCG assigns slot 1 to agent 1 at price $\epsilon - \epsilon/(r+1) + \epsilon/(r+1) - \epsilon/(2r+2)$, and slots 2 and 3 to agents 2 and 3 at prices $\epsilon/(r+1) - \epsilon/(2r+2)$ and zero. This yields revenue ϵ. In the truthful equilibrium of the VCG mechanism, on the other hand, the price is $(r+1) - (r+1)/(r+1) + \epsilon/(r+1) - \epsilon/(2r+2)$ for the first slot, $\epsilon/(r+1) - \epsilon/(2r+2)$ for the second slot, and zero for the third slot, for an overall revenue of $r + \epsilon/(r+2)$.

For α-GSP, let $v_i(\theta_i) = r+1$ for all $i \in N$. Let $\delta = \epsilon/(r+2)$, $\alpha_1 = 1$, $\alpha_2 = (1+\delta)/(r+1)$, and $\alpha_3 = 1/(r+1)$. It is easily verified that the bids $b_1 = r+1$ and $b_2 = b_3 = \delta/(1+\delta) \cdot (r+1)$ form an equilibrium of α-GSP. Given these bids, α-GSP assigns slot 1 to agent 1 at price $\delta/(1+\delta) \cdot (r+1)$, and slots 2 and 3 to agents 2 and 3 at prices δ and zero. This yields revenue $\delta/(1+\delta) \cdot (r+1) + \delta \leq \epsilon$. In the truthful equilibrium of the VCG mechanism, on the other hand, the price is r for the first slot, δ for the second slot, and zero for the third slot, for an overall revenue of $r + \delta$. □

Assuming that the click-through rate vector α is known, both α-VCG and α-GSP look very appealing: they eliminate all zero-revenue equilibria, without affecting the truthful equilibrium and without introducing any new equilibria.

In practice, however, the relevant click-through rates may not be known. A somewhat more realistic model assumes a certain degree of heterogeneity among the population generating the clicks. More precisely, a certain fraction of this population is assumed to be "merely curious," such that clicks by this part of the population do not generate any value for the agents. This introduces an information asymmetry, where the mechanism observes the overall click-through rate vector α, while agents derive value from a different click-through rate vector β.[8] In the following we will assume that β is the same for all agents, and that β is again normalized such that $\beta_1 = 1$ and $\beta_j = 0$ if $j > k$.

For a slight variation of our model, in which there is a dependence between α and β,[9] Milgrom established a separation between GSP and VCG: Suppose that the Vickrey price-per-click sequence $\{p_j(\theta)/\alpha_j\}_{j=1,...,k}$, where $p_j(\theta)$ is the truthful VCG price for slot j, is decreasing. Then α-GSP retains the VCG outcome while α-VCG fails to do so.

We obtain an analogous observation in our model.

Proposition 5. *Let $\alpha, \beta \in \mathbb{R}^n_>$. Then, α-GSP is Vickrey-preserving on Θ^β if and only if the Vickrey price-per-click sequence is decreasing.*

Proof. Fix $\alpha, \beta \in \mathbb{R}^n_>$ and $\theta \in \Theta^\beta$. Order the agents so that $v_1(\theta_1) \geq \cdots \geq v_n(\theta_n)$. Efficiency requires that agent i win position i. To get the Vickrey prices $p_i(\theta)$ for all slots i it is necessary that the equilibrium bid by agent $i \in \{2, \ldots, \min(n, k+1)\}$ be $b_i(\theta) = p_{i-1}(\theta)/\alpha_{i-1}$. Take $b_1(\theta) = \alpha_1 \cdot v_1(\theta_1)$ and $b_i(\theta) = 0$ for $i > \min(n, k+1)$.

First suppose that the sequence $\{p_j/\alpha_j\}_j$ is not decreasing, e.g., because $p_i(\theta)/\alpha_i < p_{i+1}(\theta)/\alpha_{i+1}$. It follows that $b_{i+1}(\theta) = p_i(\theta)/\alpha_i < p_{i+1}(\theta)/\alpha_i + 1 = b_{i+2}(\theta)$ and, thus, the bids are not ranked in the order required for efficiency.

Next suppose that the sequence $\{p_j/\alpha_j\}_j$ is decreasing. In this case the bids are ranked in the correct order. For a contradiction suppose that some agent j could strictly benefit from a deviation to $\theta' = (\theta_1, \ldots, \theta_{j-1}, \theta'_j, \theta_{j+1}, \ldots, \theta_n)$. Suppose that given θ' agent j is assigned slot l at price $p_l(\theta')$. Since agent j strictly benefits from the deviation we must have that $\beta_l \cdot v_j(\theta_j) - p_l(\theta') > \beta_j \cdot v_j(\theta) - p_j(\theta)$. If $l < j$, then the price that agent j faces for slot l is at least $p_l(\theta') \geq \alpha_l \cdot p_{l-1}(\theta)/\alpha_{l-1} > \alpha_l \cdot p_l(\theta)/\alpha_l = p_l(\theta)$. If $l > j$, then the price that agent j faces for slot l is exactly $p_l(\theta') = p_l(\theta)$. We conclude that $\beta_l \cdot v_j(\theta_j) - p_l(\theta) \geq \beta_l \cdot v_j(\theta_j) - p_l(\theta') > \beta_j \cdot v_j(\theta_j) - p_j(\theta)$. But this contradicts the envy freeness of the VCG assignment and payments (see, e.g., [22]). □

[8] As Milgrom [23] points out, "if the search company observes clicks but not sales or value for each position, its auction rule can entail adjusting bids in proportion to clicks but not in proportion to value" (p. 68).

[9] Milgrom assumes that there is a fraction λ of shoppers with click-through rate vector α and a fraction $(1-\lambda)$ of curious searchers with click-through rate vector β. The click-through rate vector γ observed by the search provider is then given by $\gamma_j = \lambda \cdot \alpha_j + (1-\lambda) \cdot \beta_j$.

Proposition 6. *There exist $n, k \geq 3$ and $\alpha, \beta \in \mathbb{R}^n_>$ such that the Vickrey price-per-click sequence is decreasing, but α-VCG is not Vickrey-preserving.*

Proof. Let $\alpha \in \mathbb{R}^n_>$ be such that $\alpha_1 = 1$, $\alpha_2 = 0.5$, and $\alpha_3 = 0.4$, and let $\beta \in \mathbb{R}^n_>$ be such that $\beta_1 = 1$, $\beta_2 = 0.9$, and $\beta_3 = 0.8$. Let $\theta \in \Theta^\beta$ be such that $v_1(\theta_1) = 30$, $v_2(\theta_2) = 20$, $v_3(\theta_3) = 10$, and $v_i(\theta_i) = 0$ if $i > 3$. In the VCG outcome, slot 1 is assigned to agent 1 at price $p_1(\theta) = (\beta_1 - \beta_2) \cdot v_2(\theta_2) + (\beta_2 - \beta_3) \cdot v_3(\theta_3) = 3$, slot 2 to agent 2 at $p_2(\theta) = (\beta_2 - \beta_3) \cdot v_3(\theta_3) = 1$, and slot 3 to agent 3 at $p_3(\theta) = 0$. The Vickrey price-per-click sequence $\{p_j/\alpha_j\}_j = \{3, 2, 0\}$ is decreasing. Now assume that the same outcome is obtained in α-VCG, and denote by $b \in \mathbb{R}^n$ a bid profile that leads to this outcome. Since both VCG and α-VCG are efficient, it must hold that $b_1 \geq b_2 \geq b_3$. To get the same prices as in the VCG outcome, we must further have that $b_2 = (\beta_1 - \beta_2)/(\alpha_1 - \alpha_2) \cdot v_2(\theta_2) = 4$ and $b_3 = (\beta_2 - \beta_3)/(\alpha_2 - \alpha_3) \cdot v_3(\theta_3) = 10$. Thus, $b_2 < b_3$, which yields a contradiction. □

This line of reasoning seems a bit problematic for two reasons. First, there seems no reason to believe that the condition relating the Vickrey prices and click-through rates would be satisfied in practice. Second, the above discussion only shows superiority of GSP over VCG with respect to a *particular* simplification, and it might well be the case that there exists a different simplification of the VCG mechanism with comparable or even better properties.

2.4.3 A Sense in which GSP is Superior to VCG

The above observations raise the following prominent question: *does there exist a simplification that preserves the VCG outcome despite ignorance about the true click-through rates that affect bidders' values, and if so, can this simplification achieve improved revenue relative to the VCG outcome, in every equilibrium?*

For GSP the answer is surprisingly simple: a closer look at the proofs of Proposition 4 and Proposition 5 reveals that by ignoring the observed click-through rates α, and setting $\alpha = \mathbf{1} = (1, \ldots, 1)$ instead, one obtains a simplification that is tight, guarantees positive revenue, and is Vickrey-preserving on *all* of Ψ. This strengthens Proposition 4 over the claims for α-GSP and α-VCG.

Corollary 1. *α-GSP is tight on Ψ, has positive revenue on Ψ if $n, k \geq 2$, and is Vickrey-preserving on Ψ, if and only if $\alpha = \mathbf{1} = (1, \ldots, 1)$.*

The direction from left to right follows by observing that, for every $\alpha \neq \mathbf{1}$, we can find a β such that the condition of Proposition 5 is violated.

In light of Theorem 1, and given the arguments in favor of efficient equilibria, we may further ask for the minimum revenue obtained by $\mathbf{1}$-GSP in any efficient equilibrium. It turns out that $\mathbf{1}$-GSP always recovers at least half of the VCG revenue for all slots but the first.

Theorem 2. *Let $\beta \in \mathbb{R}^n_>$, $\theta \in \Theta^\beta$, and assume that $v_1(\theta_1) \geq \cdots \geq v_n(\theta_n)$. Then, every efficient equilibrium of* **1***-GSP for* θ *yields revenue at least*

$$\frac{1}{2}\left(R^{VCG}(\theta) - \sum_{j=1}^{k}(\beta_j - \beta_{j+1}) \cdot v_{j+1}(\theta_{j+1})\right).$$

Proof. Let $b(\theta)$ be a bid profile corresponding to an efficient equilibrium of **1**-GSP. It then holds that $b_1(\theta) \geq \cdots \geq b_n(\theta)$, and for all $i \in \{1,\ldots,k\}$, **1**-GSP assigns slot i to agent i at price $p_i(\theta) = b_{i+1}(\theta)$. A necessary condition for $b(\theta)$ to be an equilibrium is that for every agent $j \in N$, $b_j(\theta)$ is large enough such that none of the agents $i > j$ would prefer being assigned slot j instead of i. In particular, for every $i \in N$,

$$\beta_{i+1} \cdot v_{i+1}(\theta_{i+1}) - p_{i+1}(\theta) \geq \beta_i \cdot v_{i+1}(\theta_{i+1}) - b_i(\theta) \quad \text{and}$$
$$\beta_{i+2} \cdot v_{i+2}(\theta_{i+2}) - p_{i+2}(\theta) \geq \beta_i \cdot v_{i+2}(\theta_{i+2}) - b_i(\theta).$$

Since $p_i(\theta) = b_{i+1}(\theta)$ and by rearranging,

$$b_i(\theta) \geq (\beta_i - \beta_{i+1}) \cdot v_{i+1}(\theta_{i+1}) + p_{i+1}(\theta) = (\beta_i - \beta_{i+1}) \cdot v_{i+1}(\theta_{i+1}) + b_{i+2}(\theta) \quad \text{and}$$
$$b_i(\theta) \geq (\beta_i - \beta_{i+2}) \cdot v_{i+2}(\theta_{i+2}) + p_{i+2}(\theta) \geq (\beta_{i+1} - \beta_{i+2}) \cdot v_{i+2}(\theta_{i+2}) + b_{i+3}(\theta).$$

If we repeatedly substitute according to the first inequality, we obtain

$$b_i(\theta) \geq \sum_{j=1}^{\lfloor \frac{k-i}{2} \rfloor}(\beta_{i+2\cdot j-2} - \beta_{i+2\cdot j-1}) \cdot v_{i+2\cdot j-1}(\theta_{i+2\cdot j-1}) \quad \text{and}$$
$$b_i(\theta) \geq \sum_{j=1}^{\lfloor \frac{k-i}{2} \rfloor}(\beta_{i+2\cdot j-1} - \beta_{i+2\cdot j}) \cdot v_{i+2\cdot j}(\theta_{i+2\cdot j}).$$

By adding the two inequalities,

$$2 \cdot b_i(\theta) \geq \sum_{j=i}^{k}(\beta_j - \beta_{j+1}) \cdot v_{j+1}(\theta_{j+1}),$$

and, since $p_i(\theta) = b_{i+1}(\theta)$,

$$2 \cdot p_i(\theta) \geq \sum_{j=i+1}^{k}(\beta_j - \beta_{j+1}) \cdot v_{j+1}(\theta_{j+1}).$$

Now recall that $R^{VCG}(\theta) = \sum_{i \in N} r_i(\theta)$, where

$$r_i(\theta) = \sum_{j=i}^{k}(\beta_j - \beta_{j+1}) \cdot v_{j+1}(\theta_{j+1}).$$

Thus,

$$\sum_{i \in N} 2 \cdot p_i(\theta) \geq \sum_{i \in N} r_i(\theta) - \sum_{j=1}^{k}(\beta_j - \beta_{j+1}) \cdot v_{j+1}(\theta_{j+1}).$$

The revenue obtained in any efficient equilibrium of **1**-GSP is therefore at least

$$\sum_{i \in N} p_i(\theta) \geq \frac{1}{2}\left(R^{VCG}(\theta) - \sum_{j=1}^{k} (\beta_j - \beta_{j+1}) \cdot v_{j+1}(\theta_{j+1})\right).$$

□

Our analysis also leads to a satisfactory contrast between **1**-GSP and simplifications of VCG: *any* simplification of VCG that does not observe the value-generating click-through rates, and is Vickrey-preserving for all possible choices of these click-through rates, must admit an efficient equilibrium with arbitrarily low revenue.

Theorem 3. *Let $\hat{M}$ be a simplification of the VCG mechanism that is Vickrey-preserving on Ψ. Then, for every $\theta \in \Psi$ and every $\epsilon > 0$, there exists an efficient equilibrium of $\hat{M}$ with revenue at most ϵ.*

Proof. Fix $\beta \in \mathbb{R}^n_>$ and consider an arbitrary type profile $\theta \in \Theta^\beta \subseteq \Psi$. Order the agents such that $v_1(\theta_1) \geq \cdots \geq v_n(\theta_n)$.

The proof proceeds in two steps. First we will argue that for some $c \geq 0$, every $\delta > 0$, and all $i \in N$, $\hat{X}_i$ must contain a message x_i^δ corresponding to bids b_{ij} such that $b_{ii} = \beta_i \cdot v_i(\theta_i) + c$, $b_{ij} \leq \beta_i \cdot v_i(\theta_i) + c + \delta$ for $1 \leq j < i$ and $b_{ij} \leq c + \delta$ for $i < j \leq k$. These messages will then be used to construct an equilibrium with low revenue.

To show that the restricted message spaces $\hat{X}_i$ must contain messages as described above, we show that these messages are required to reach the VCG outcome for a different type profile $\theta' \in \Theta^{\beta'} \subseteq \Psi$ for a particular $\beta' \in \mathbb{R}^n_>$. Denote by $p_i(\theta')$ the price of slot i for type profile θ'. We know that for $j \in \{2, \ldots, k\}$, $\beta'_j \cdot v_j(\theta'_j) = \beta'_{j-1} \cdot v_j(\theta'_j) - p_{j-1}(\theta')$ (see, e.g., [23]). Consider an arbitrary $\delta > 0$. If we choose β' such that $\beta'_1 - \beta'_i$ and β'_{i+1} are small enough, we can choose θ' as above such that

$$\begin{aligned} &p_j(\theta') - p_i(\theta') \leq \delta && \text{for } j < i \text{ and} \\ &p_i(\theta') - p_j(\theta') \geq \beta_i \cdot v_i(\theta_i) - \delta && \text{for } j > i. \end{aligned}$$

A well-known property of the VCG outcome in the assignment problem is its envy freeness (see, e.g., [22]): denoting by b_{ij} the bid of agent i on slot j and by p_j the price of slot j, it must hold for every agent i that

$$b_{ii} - p_i \geq b_{ij} - p_j \qquad \text{for all } j \in S.$$

For type profile θ', we thus obtain

$$\begin{aligned} &b_{ij} - b_{ii} \leq p_j(\theta') - p_i(\theta') \leq \delta && \text{for } j < i \text{ and} \\ &b_{ii} - b_{ij} \geq p_i(\theta') - p_j(\theta') \geq \beta_i \cdot v_i(\theta_i) - \delta && \text{for } j > i. \end{aligned}$$

Using messages x_i^δ, we now construct an efficient equilibrium with low revenue. Clearly, the allocation that assigns slot i to agent i is still efficient under message profile x^δ. Furthermore, for all $j \in \{1,\ldots,k\}$, the VCG price of slot j under x^δ goes to zero as δ goes to zero. In particular, we can choose δ such that the overall revenue is at most ϵ. We finally claim that there exists some $\delta' > 0$ such that x^δ is an equilibrium for every δ with $\delta' > \delta > 0$. To see this, recall that $\beta_i \cdot v_i(\theta_i) > \beta_j \cdot v_i(\theta_i)$ for $j > i$, so $u_i(x^\delta,\theta_i) > \beta_j \cdot v_i(\theta_i)$ for some small enough δ. If, on the other hand, agent i was assigned a slot $j < i$, his payment would be at least $\beta_j \cdot v_j(\theta_j) - \delta > \beta_j \cdot v_i(\theta_i) - \delta$. This would leave him with utility at most δ, which can be chosen to be smaller than $u_i(x^\delta,\theta_i)$. □

It is worth noting that despite having a reasonably good lower bound on revenue, **1**-GSP does not quite succeed in circumventing Theorem 1: there exists a type profile for which only the first slot generates a significant amount of VCG revenue, and an equilibrium of **1**-GSP for this type profile in which revenue is close to zero. On the other hand, for a wide range of click-through rates, there will be a large gap between the minimum revenue of **1**-GSP and simplifications of VCG in efficient equilibria.

Moreover, the revenue separation between simplifications of GSP and VCG only applies to simplifications of the latter that are Vickrey-preserving. Given how central the VCG outcome has been to the analysis of sponsored search auctions, this seems a reasonable property to impose. Nevertheless, it is an interesting question whether the result can be strengthened further.

2.5 Combinatorial Auctions

Mechanisms for combinatorial auctions allocate items from a set G to the agents, i.e., $\Omega = \prod_{i\in N} 2^G$ such that for every $o \in \Omega$ and $i, j \in N$ with $i \neq j$, $o_i \cap o_j = \emptyset$. We make the standard assumption that the empty set is valued at zero and that valuations satisfy free disposal, i.e., for all $i \in N$, $\theta_i \in \Theta_i$, and $o, o' \in \Omega$, $v_i(o,\theta_i) = 0$ if $o_i = \emptyset$ and $v_i(o,\theta_i) \leq v_i(o',\theta_i)$ if $o_i \subseteq o'_i$. The latter condition also implies that each agent is only interested in the package he receives, and we sometimes abuse notation and write $v_i(C,\theta_i)$ for the valuation of agent i for any $o \in \Omega$ with $o_i = C$. We further write $k = |G|$ for the number of items, $W(o,x) = \sum_{j\in N} v_j(o,x_j)$ for the social welfare of outcome $o \in \Omega$ under message profile $x \in X$, and $W_{max}(x) = \max_{o\in\Omega} W(o,x)$ for the maximum social welfare of any outcome. Finally, for every agent $i \in N$, message $x_i \in X$, and bundle of items $B \subseteq G$ we write $x_i(B)$ for agent i's bid on bundle B.

The VCG mechanism makes it a dominant strategy for every agent to bid his true valuation for every bundle of items. Since the number of such bundles is exponential in the number of items, however, computational constraints might prevent agents from playing this dominant strategy (even for a well-crafted bidding language [25, 20]). In light of these results, and in light of the observation that simplifications can help to isolate useful equilibria, it is interesting to ask which other (ex-post) equilibria the VCG auction can have. Holzman and Monderer [16] showed that these equilibria are precisely the projections of the true types to those subsets of the set of all bundles that form a quasi-field. Let $\Sigma \subseteq 2^G$ be a set of bundles of items such that $\emptyset \in \Sigma$. Σ is called a *quasi-field* if it is closed under complementation and union of disjoint subsets, i.e., if

- $B \in \Sigma$ implies $B^c \in \Sigma$, where $B^c = G \setminus B$ and
- $B, C \in \Sigma$ and $B \cap C$ implies $B \cup C \in \Sigma$.

For a message $x_i \in X_i$, write x_i^Σ for the projection of x_i to Σ, i.e., for the unique message such that for every bundle $B \subseteq G$ of items,

$$x_i^\Sigma(B) = \max_{B' \in \Sigma, B' \subseteq B} x_i(B').$$

The characterization given by Holzman and Monderer is subject to the additional constraint of *variable participation*: a strategy profile s for a set N of agents is an equilibrium of a VCG mechanism under variable participation if for every $N' \subseteq N$, the projection of s to N' is an equilibrium of every VCG mechanism for N'.

Theorem 4 (Holzman and Monderer [16])**.** *Consider a VCG combinatorial auction with a set N of agents and a set G of items. Then, a strategy profile $s = (s_1, \ldots, s_n)$ is an ex-post equilibrium of this auction under variable participation if and only if there exists a quasi-field $\Sigma \subseteq 2^G$ such that for every type profile θ and every agent $i \in N$, $s_i(\theta_i) = \theta_i^\Sigma$.*

Intuitively, the social welfare obtained in these "bundling" equilibria decreases as the set of bundles becomes smaller. A simple argument shows, for example, that welfare in the bundling equilibrium for Σ, where $|\Sigma| = 2^m$ for some $m \leq k$, can be smaller by a factor of k/m than the maximum welfare. For this, consider a setting with k agents such that each agent desires exactly one of the items, i.e., values this item at 1, and each item is desired by exactly one of the agents. Clearly, maximum social welfare in this case is k. On the other hand, since Σ is a quasi-field, it cannot contain more than m bundles that are pairwise disjoint. Therefore, by assigning only bundles in Σ, one can obtain welfare at most m.

Welfare can also differ tremendously among quasi-fields of equal size, which suggests an opportunity for simplification.

Proposition 7. *Let G be a set of items, $k = |G|$, and $m \leq k$. Then there exist quasi-fields $\Sigma, \Sigma' \subseteq 2^G$ with $|\Sigma| = |\Sigma'| = 2^m$ and a type profile θ such that*

$$\frac{W_{max}(\theta^\Sigma)}{W_{max}(\theta^{\Sigma'})} \geq \frac{m}{\lceil m^2/k \rceil}.$$

Proof. Consider a partition of G into sets $G_1, \ldots, G_m$ of size $\lceil \frac{k}{m} \rceil$ or $\lfloor \frac{k}{m} \rfloor$, and let Σ be the closure of $\{G_1, \ldots, G_m\}$ under complementation and union of disjoint sets. For every i, $1 \leq i \leq m$, choose an arbitrary $g_i \in G_i$, and define $\theta = (\theta_1, \ldots, \theta_m)$ such that for every i with $1 \le i \le m$, $v_i(X, \theta_i) = 1$ if $g_i \in X$ and $v_i(X, \theta_i) = 0$ otherwise. Clearly, $W_{max}(\theta^{\Sigma}) = m$. Now consider a second partition of G into sets $G'_1, \ldots, G'_m$ of size $\lceil \frac{k}{m} \rceil$ or $\lfloor \frac{k}{m} \rfloor$ such that $G'_j \supseteq \{g_i : (j-1)\lfloor k/m \rfloor + 1 \le i \le j\lfloor k/m \rfloor\}$, and let Σ' be the closure of $\{G_1, \ldots, G_m\}$ under complementation and union of disjoint sets. It is then easily verified that $W_{max}(\theta^{\Sigma'}) \le \lceil m/(k/m) \rceil = \lceil m^2/k \rceil$, and the claim follows. □

Agents might also disagree about the quality of the different bundling equilibria of a given maximum size. In particular, the set of these equilibria might contain several Pareto undominated equilibria, but no dominant strategy equilibrium. From the point of view of equilibrium selection, this is the worst possible scenario.

Example 1. Let $N = \{1, 2, 3\}$, $G = \{A, B, C\}$, and consider a type profile $\theta = (\theta_1, \theta_2, \theta_3)$ such that

$$v_1(X, \theta_1) = \begin{cases} 4 & \text{if } \{A, C\} \subseteq X, \\ 3 & \text{if } A \in X \text{ and } C \notin X \\ 1 & \text{otherwise;} \end{cases}$$

$$v_2(X, \theta_2) = \begin{cases} 3 & \text{if } \{A, C\} \subseteq X \text{ or } \{B, C\} \subseteq X \text{ and} \\ 0 & \text{otherwise;} \end{cases}$$

$$v_3(X, \theta_3) = \begin{cases} 1 & \text{if } B \in X \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, at least four bids are required to express θ_1. Since a quasi-field on G must contain both the empty set and G itself, there are four quasi-fields of size four or less:

$$\Sigma^1 = \{\emptyset, \{A\}, \{B, C\}, \{A, B, C\}\} \qquad \Sigma^2 = \{\emptyset, \{B\}, \{A, C\}, \{A, B, C\}\}$$
$$\Sigma^3 = \{\emptyset, \{C\}, \{A, B\}, \{A, B, C\}\} \qquad \Sigma^4 = \{\emptyset, \{A, B, C\}\}$$

For $i \in \{1, 2, 3, 4\}$, write $\theta^i = (\theta_1^{\Sigma^i}, \theta_2^{\Sigma^i}, \theta_3^{\Sigma^i})$ for the projection of θ to Σ^i. The following is now easily verified. In the VCG outcome for θ^1, agent 1 is assigned $\{A\}$ at price 0 for a utility of 3, and agent 2 is assigned $\{B, C\}$ at price 1 for a utility of 2. In the VCG outcome for θ^2, agent 1 is assigned $\{A, C\}$ at price 3 for a utility of 1, and agent 3 is assigned $\{B\}$ at price 0 for a utility of 1. Finally, in the VCG outcome for θ^3 and θ^4, agent 1 is assigned $\{A, B, C\}$ at price 3 for a utility of 1. The outcomes for θ^1 and θ^2 are both Pareto undominated. Also observe that social welfare is greater for θ^1, while θ^2 yields higher revenue.

Finally, the projection to a quasi-field can result in an equilibrium with revenue zero, even if revenue in the dominant strategy equilibrium is strictly positive. This is illustrated in the following example. It should be noted that this example, as well as the previous one, can easily be generalized to arbitrary numbers of agents and items and a large range of upper bounds on the size of the quasi-field.

Example 2. Let $N = \{1,2,3\}$, $G = \{A,B,C,D\}$, and consider a type profile $\theta = (\theta_1, \theta_2, \theta_3)$ such that

$$v_1(X,\theta_1) = \begin{cases} 2 & \text{if } \{A,D\} \subseteq X \text{ and} \\ 0 & \text{otherwise;} \end{cases}$$

$$v_2(X,\theta_2) = \begin{cases} 2 & \text{if } \{A,B\} \subseteq X, \\ 1 & \text{if } B \in X \text{ and } A \notin X, \text{ and} \\ 0 & \text{otherwise;} \end{cases}$$

$$v_3(X,\theta_3) = \begin{cases} 2 & \text{if } C \in X, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

In the VCG outcome for θ, agent 1 is assigned $\{A,D\}$ at price 1, agent 2 is assigned $\{B\}$ at price 0, and agent 3 is assigned $\{C\}$ at price 0, for an overall revenue of 1. In the VCG outcome for θ^Σ, on the other hand, where $\Sigma = \{\emptyset, AB, CD, ABCD\}$, agent 2 is assigned $\{A,B\}$ and agent 3 is assigned $\{C,D\}$, both at price 0. Revenue is 0 as well.

In theory, one way to solve these problems is to simplify the mechanism, and artificially restrict the set of bundles agents can bid on. Given a set $\Sigma \subseteq 2^G$ of bundles, call Σ-VCG the simplification of the VCG mechanism obtained by restricting the message spaces to $\hat{X}_i \subseteq X_i$ such that for every $\hat{x}_i \in \hat{X}_i$ and every bundle of items $B \subseteq G$,

$$\hat{x}_i(B) = \max_{B' \in \Sigma, B' \subseteq B} \hat{x}_i(B').$$

In other words, Σ-VCG allows agents to bid only on elements of Σ and derives bids for the other bundles as the maximum bid on a contained bundle. It is easy to see that Σ-VCG is *maximal in range* [24], i.e., that it maximizes social welfare over a subset of Ω. It follows that for each agent, truthful projection onto Σ is a dominant strategy in Σ-VCG.

This shows that simplification can focus attention on a focal (truthful) equilibrium and thus avoid equilibrium selection along a Pareto frontier. As Proposition 7 and the above examples suggest, this can have a significant positive impact on both social welfare and revenue. It does not tell us, of course, how Σ should be chosen in practice. One understood approach for maximizing social welfare without any knowledge about the quality of different outcomes, and without consideration to Σ being a quasi-field, is to partition G arbitrarily into m sets of roughly equal size, where m is the largest number of bundles agents can bid on. The welfare thus obtained is smaller than the maximum social welfare by a factor of at most $k/\sqrt{m}$ [17]. If additional knowledge is available, however, it may be possible to improve the result substantially, as the above results comparing the outcomes for different values of Σ show.

With that being said, there are (at least) two additional properties that are desirable for a simplification: that Σ-VCG is tight, and that Σ is a quasi-field. Tightness ensures that no additional equilibria are introduced as compared to the fully expressive VCG mechanism, such that the quality of the worst equilibrium outcome of the simplification is no worse than that of the original mechanism. This remains important when Σ may still be too large for agents to use its full projection, in which case agents would again have to select from a large set of possible ex-post equilibria. By requiring that Σ is a quasi-field, in addition to being a dominant strategy equilibrium of Σ-VCG, truthful projection to Σ is also an ex-post equilibrium of the fully expressive VCG mechanism, and thus stable against unrestricted unilateral deviations. This ensures that agents do not experience regret, in the sense of being prevented from sending a message they would want to send given the messages sent by the other agents.

It turns out that these two requirements are actually equivalent, i.e., that Σ-VCG is tight if and only if Σ is a quasi-field. The following result holds with respect to both Nash equilibria and ex-post equilibria.

Theorem 5. *Let $\Sigma \subseteq 2^G$ such that $\emptyset \in \Sigma$. Then, Σ-VCG is a tight simplification if Σ is a quasi-field, and this condition is also necessary if $n \geq 3$.*

Proof. For the direction from right to left, assume that Σ is a quasi-field. By Proposition 1 it suffices to show that Σ-VCG satisfies outcome closure. Fix valuation functions v_j and types θ_j for all $j \in N$, and consider an arbitrary agent $i \in N$. We claim that for every $x_i \in X_i$ and every $\hat{x}_{-i} \in \hat{X}_{-i}$,

$$u_i((\theta_i^\Sigma, \hat{x}_{-i}), \theta_i) \geq u_i((x_i, \hat{x}_{-i}), \theta_i).$$

To see this, observe that there exists $\hat{\theta}_{-i} \in \Theta_{-i}$ such that $\hat{\theta}^\Sigma_{-i} = \hat{x}_{-i}$, and consider the type profile $(\theta_i, \hat{\theta}_{-i})$. Holzman et al. [17] have shown that the projection of the true types to a quasi-field Σ is an ex-post equilibrium of the (fully expressive) VCG mechanism. Thus, in particular, θ_i^Σ is a best response to $\hat{\theta}^\Sigma_{-i} = \hat{x}_{-i}$, which proves the claim.

For the direction from left to right, assume that Σ is not a quasi-field. Holzman et al. [17] have shown that in this case, θ^Σ is not an ex-post equilibrium of the VCG mechanism. On the other hand, θ^Σ is a dominant-strategy equilibrium of Σ-VCG, because Σ-VCG is maximal in range. This shows that Σ-VCG is not tight. □

Together with Theorem 4, this yields a characterization of the ex-post equilibria of Σ-VCG for the case when Σ is a quasi-field.

Corollary 2. *Let Σ be a quasi-field. Then, $\hat{x}$ is an ex-post equilibrium of Σ-VCG under variable participation if and only if $\hat{x} = \theta^{\Sigma'}$ for some quasi-field $\Sigma' \subseteq \Sigma$.*

Since a quasi-field of size m can contain at most $\log m$ bundles that are pairwise disjoint, insisting that a simplification be tight does come at a cost, decreasing the worst-case social welfare in the truthful projection by an additional factor of up to $\sqrt{m/\log m}$. Still, as discussed above, tightness brings other advantages to the simplified mechanism.

2.6 The Role of Information

Like most of the literature on sponsored search auctions, we have assumed that agents have complete information about each others' valuations for the different slots. It turns out that this assumption is crucial, and that the positive results for α-GSP do not extend to incomplete information settings.

In particular, α-GSP admits an ex-post equilibrium only in very degenerate cases. This complements a result of Gomes and Sweeney [14], who showed that α-GSP has an efficient Bayes-Nash equilibrium on type space Θ^α if and only if α decreases sufficiently quickly. Of course, ex-post equilibrium is stronger than Bayes-Nash equilibrium, but on the other hand our result precludes the existence of the former for a significantly larger type space.

Theorem 6. *Let $\alpha, \beta \in \mathbb{R}^n_>$. Then, α-GSP has an efficient ex-post equilibrium on type space Θ^β if and only if $n \leq 2$ or $k \leq 1$, even if $\alpha = \beta$.*

Proof. If $n = 1$ or $k = 1$, strategies s_i with $s_i(\theta_i) = \beta_1/\alpha_1 \cdot v_i(\theta_i)$ ensure that the agent with the highest valuation is assigned the first slot at a price equal to the second-highest valuation. It is easy to see that this is efficient and constitutes an ex-post equilibrium.

Now consider the case where $n = 2$ and $k \geq 2$. We claim that $s_i(\theta_i) = (\beta_1 - \beta_2)/\alpha_1 \cdot v_i(\theta_i)$ for $i \in \{1,2\}$ is the unique efficient ex-post equilibrium in this case. To see this, observe that s_1 is an equilibrium strategy if and only if for all $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$,

$$\begin{aligned} &\beta_1 \cdot v_2(\theta_2) - \alpha_1 \cdot s_1(\theta_1) \geq \beta_2 \cdot v_2(\theta_2) \quad \text{if agent 2 is assigned the first slot, and} \\ &\beta_1 \cdot v_2(\theta_2) - \alpha_1 \cdot s_1(\theta_1) \leq \beta_2 \cdot v_2(\theta_2) \quad \text{if agent 2 is assigned the second slot.} \end{aligned}$$

For s_1 to be part of an efficient equilibrium, the first inequality has to hold if $v_2(\theta_2) = v_1(\theta_1) + \epsilon$ for any $\epsilon > 0$, and the second inequality has to hold if $v_2(\theta_2) = v_1(\theta_1) - \epsilon$ for any $\epsilon > 0$. By rearranging, we get that for every $\theta_1 \in \Theta_1$ and every $\epsilon > 0$,

$$s_1(\theta_1) \leq \frac{\beta_1 - \beta_2}{\alpha_1} \cdot (v_1(\theta_1) + \epsilon) \qquad \text{and} \qquad s_1(\theta_1) \geq \frac{\beta_1 - \beta_2}{\alpha_1} \cdot (v_1(\theta_1) - \epsilon).$$

Since analogous conditions have to hold for s_2, the claim follows.

Finally consider the case where $n \geq 3$ and $k \geq 2$, and assume for contradiction that s is an efficient ex-post equilibrium. Observe that s must remain an equilibrium if we restrict the types in such a way that $v_\ell(\theta_\ell) > 0$ if $\ell \in \{1,2\}$ and $v_\ell(\theta_\ell) = 0$ otherwise. Strategies s_1 and s_2 thus have to be of the form described above, i.e., $s_i(\theta_i) = (\beta_1 - \beta_2)/\alpha_1 \cdot v_i(\theta_i)$ for $i \in \{1,2\}$. Similarly, s remains an equilibrium if we restrict the valuations such that $v_\ell(\theta_\ell) > 0$ if $\ell = 3$

and $v_\ell(\theta_\ell) = 0$ otherwise. It follows that $s_3(\theta_3) > 0$ if $v_3(\theta_3) > 0$. Let $v_1(\theta_1) = v_2(\theta_2) = v$ for some $v > 0$, choose θ_3 such that $0 < v_3(\theta_3) < (\beta_1 - \beta_2)/\alpha_1 \cdot v$, and let $v_\ell(\theta_\ell) = 0$ for $\ell > 3$. Thus $s_1(\theta_1) = (\beta_1 - \beta_2)/\alpha_1 \cdot v$, $s_2(\theta_2) = (\beta_1 - \beta_2)/\alpha_1 \cdot v$, and $s_\ell(\theta_\ell) = 0$ for $\ell > 3$. It further holds that $v_3(\theta_3) < v$, because $\beta_1 > \beta_2$ and $\alpha_1 = 1$. We may now assume without loss of generality that for this bid profile, α-GSP assigns slot 1 to agent 1, slot 2 to agent 2, and slot 3 to agent 3; the case where slot 1 is assigned to agent 2 and slot 2 to agent 1 is symmetric. Agent 2 thus obtains utility $u_2 = \beta_2 \cdot v - s_3(\theta_3) < \beta_2 \cdot v$. If he changed his bid to $b_2' > b_2$, he would be assigned slot 1 at price $p_1' = \alpha_1 \cdot (\beta_1 - \beta_2)/\alpha_1 \cdot v = (\beta_1 - \beta_2) \cdot v$, and obtain utility $u_2' = \beta_1 \cdot v - p_1' = \beta_1 \cdot v - (\beta_1 - \beta_2) \cdot v = \beta_2 \cdot v > u_2$. This contradicts the assumption that s is an equilibrium. □

The α-VCG mechanism does only slightly better: it has an ex-post equilibrium only when it allows agents to bid truthfully (it is thus not a meaningful simplification, in that the language is in this sense exact).

Theorem 7. *Let $\alpha, \beta \in \mathbb{R}^n_>$. Then, α-VCG has an efficient ex-post equilibrium on type space Θ^β if and only if $\alpha = \beta$, or $n \le 2$, or $k \le 1$.*

Proof. If $\alpha = \beta$, agent i can bid truthfully on all slots by letting $s_i(\theta_i) = v_i(\theta_i)$. If $n = 1$ or $k = 1$, only one slot will be assigned, and agent i can bid truthfully on this slot by letting $s_i(\theta_i) = \beta_1/\alpha_1 \cdot v_i(\theta_i)$. In both cases, truthful bidding constitutes an efficient dominant strategy equilibrium and, a fortiori, an efficient ex-post equilibrium.

Now consider the case where $n = 2$ and $k \ge 2$. We claim that strategy profile s where $s_i(\theta_i) = (\beta_1 - \beta_2)/(\alpha_1 - \alpha_2) \cdot v_i(\theta_i)$ for $i \in \{1, 2\}$ is the unique efficient ex-post equilibrium in this case. To see this, observe that s_1 is an equilibrium strategy if and only if for all $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$,

$$\beta_1 \cdot v_2(\theta_2) - (\alpha_1 - \alpha_2) \cdot s_1(\theta_1) \ge \beta_2 \cdot v_2(\theta_2) \quad \text{if agent 2 is assigned the first slot, and}$$
$$\beta_1 \cdot v_2(\theta_2) - (\alpha_1 - \alpha_2) \cdot s_1(\theta_1) \le \beta_2 \cdot v_2(\theta_2) \quad \text{if agent 2 is assigned the second slot.}$$

For s_1 to be part of an efficient equilibrium, the first inequality has to hold if $v_2(\theta_2) = v_1(\theta_1) + \epsilon$ for any $\epsilon > 0$, and the second inequality has to hold if $v_2(\theta_2) = v_1(\theta_1) - \epsilon$ for any $\epsilon > 0$. By rearranging, we get that for every $\theta_1 \in \Theta_1$ and every $\epsilon > 0$,

$$s_1(\theta_1) \le \frac{\beta_1 - \beta_2}{\alpha_1 - \alpha_2} \cdot (v_1(\theta_1) + \epsilon) \qquad \text{and} \qquad s_1(\theta_1) \ge \frac{\beta_1 - \beta_2}{\alpha_1 - \alpha_2} \cdot (v_1(\theta_1) - \epsilon).$$

Since analogous conditions have to hold for s_2, the claim follows.

Finally consider the case where $n \ge 3$, $k \ge 2$, and $\alpha \ne \beta$, and assume for contradiction that s is an efficient ex-post equilibrium. We first claim that s is symmetric. For a contradiction assume that there exist agents i and j such that $s_i \ne s_j$. Observe that s must remain an equilibrium if we restrict the valuations of all other agents such that in every efficient assignment, agents i and j are assigned the last two slots. This means, however, that s_i and s_j induce an efficient ex-post equilibrium for the case where $n = 2$ and $k \geq 2$. By assumption, this equilibrium is asymmetric, contradicting an observation we have made in the previous paragraph.

Since s is symmetric, there must exist a function $s_* : \mathbb{R} \to \mathbb{R}$ such that for every $i \in N$ and every $\theta_i \in \Theta_i$, $s_i(\theta_i) = s_*(v_i(\theta_i))$. It is not hard to see that $s_*(v) > 0$ if $v > 0$ and $s_*(v) = 0$ if $v = 0$. Since by convention $\alpha_1 = \beta_1 = 1$ and $\alpha_{k+1} = \beta_{k+1} = 0$, and since $\alpha_j \neq \beta_j$ for some $j < k$, there must further exist $i > 1$ such that (i) $\alpha_j - \alpha_{j+1} \leq \beta_j - \beta_{j+1}$ for all $j < i$ and $\alpha_i - \alpha_{i+1} > \beta_i - \beta_{i+1}$, or (ii) $\alpha_j - \alpha_{j+1} \geq \beta_j - \beta_{j+1}$ for all $j < i$ and $\alpha_i - \alpha_{i+1} < \beta_i - \beta_{i+1}$. Since the two cases are symmetric, it suffices to consider the first one.

Consider an arbitrary $v > 0$. We distinguish two cases.

First assume that $s_*(v) > (\beta_i - \beta_{i+1})/(\alpha_i - \alpha_{i+1}) \cdot v$. Let $v_\ell(\theta_\ell) = v$ for $\ell \leq i+1$ and $v_\ell(\theta_\ell) = 0$ for $\ell > i+1$. It then holds that $s_\ell(\theta_\ell) = s_*(v) > 0$ for $\ell \leq i+1$ and $s_\ell(\theta_\ell) = s_*(0) = 0$ for $\ell > i+1$. Assume without loss of generality that for this bid profile, α-VCG assigns slot ℓ to agent ℓ for $\ell \leq i+1$, and that the remaining agents are not assigned a slot. The utility of agent i in this case is $u_i = \beta_i \cdot v - p_i(\theta) = \beta_i \cdot v - (\alpha_i - \alpha_{i+1}) \cdot s_*(v)$. If he changed his bid to b with $0 < b < s_*(v)$, he would be assigned slot $i+1$ at price 0 and obtain utility $\beta_{i+1} \cdot v = \beta_{i+1} \cdot v + (\alpha_i - \alpha_{i+1}) \cdot s_*(v) - p_i > \beta_i \cdot v - p_i = u_i$. This contradicts the assumption that s is an equilibrium.

Now assume that $s_*(v) \leq (\beta_i - \beta_{i+1})/(\alpha_i - \alpha_{i+1}) \cdot v$, and observe that in this case $s_*(v) < v$. Let $v_\ell(\theta_\ell) = v$ for $\ell \leq i$ and $v_\ell(\theta_\ell) = 0$ for $\ell > i$. It then holds that $s_\ell(\theta_\ell) = s_*(v) > 0$ for $\ell \leq i$ and $s_\ell(\theta_\ell) = s_*(0) = 0$ for $\ell > i$. Assume without loss of generality that for this bid profile, α-VCG assigns slot ℓ to agent ℓ for $\ell \leq i$, and that the remaining agents are not assigned a slot. The utility of agent i under this assignment is $u_i = \beta_i \cdot v - p_i(\theta) = \beta_i \cdot v$. If he changed his bid to $b > s_*(v)$, he would be assigned slot 1 at price $\sum_{j<i}(\alpha_j - \alpha_{j+1}) \cdot s_*(v)$ and obtain utility $u'_i = \beta_1 \cdot v - \sum_{j<i}(\alpha_j - \alpha_{j+1}) \cdot s_*(v) > \beta_1 \cdot v - \sum_{j<i}(\beta_j - \beta_{j+1}) \cdot v = \beta_i \cdot v = u_i$. This again contradicts the assumption that s is an equilibrium. □

These results indicate that simplification is not very useful in sponsored search auctions without the assumption of complete information amongst bidders. Interestingly, simplification is also not necessary in this case to preclude equilibria with bad properties, at least in the special case of our model where the valuations are proportional to some (possibly unknown) vector of value-generating click-through rates. In particular, payments in *every* efficient ex-post equilibrium of the fully expressive VCG mechanism equal the VCG payments. Moreover, truthful reporting is the only efficient ex-post equilibrium when the number of agents is greater than the number of slots.

Proposition 8. *Consider a VCG sponsored search auction with type profile $\theta \in \Theta^\alpha$, and assume that $s = (s_1, \ldots, s_n)$ is an efficient ex-post equilibrium. Then, for all i with $1 \leq i \leq k$, the payment for slot i in the outcome for strategy profile s equals $p_i = \sum_{j=i}^{\min(k,n-1)}((\alpha_j - \alpha_{j+1}) \cdot v_{j+1}(\theta_{j+1}))$. Moreover, if $n > k$, then $s_i(k, \theta_i) = \alpha_k \cdot v_i(\theta_i)$ for all i with $1 \leq i \leq n$.*

Proof. Consider two agents $i, i' \in N$ and two consecutive slots j and $j+1$. Fix θ_i. Since the strategy s_i of agent i does not depend on the types of the other agents, we can choose the types of the agents in $N \setminus \{i, i'\}$ in such a way that an efficient outcome assigns slot j to agents i and slot $j+1$ to agent i', or vice versa. Doing so also fixes the bids of all agents in $N \setminus \{i, i'\}$, and thus the payment p_{j+1} associated with slot $j+1$.

For s_i to be part of an ex-post equilibrium it must hold that

$$\alpha_j \cdot v_{i'}(\theta_{i'}) - (s_i(j,\theta_i) - s_i(j+1,\theta_i)) - p_{j+1} \geq \alpha_{j+1} \cdot v_{i'}(\theta_{i'}) - p_{j+1}$$

if agent i' is assigned slot j, and

$$\alpha_{j+1} \cdot v_{i'}(\theta_{i'}) - p_{j+1} \geq \alpha_j \cdot v_{i'}(\theta_{i'}) - (s_i(j,\theta_i) - s_i(j+1,\theta_i)) - p_{j+1}$$

if agent i' is assigned slot $j+1$. For the equilibrium to be efficient the first inequality must hold if $v_{i'}(\theta_{i'}) = v_i(\theta_i) + \epsilon$ for any $\epsilon > 0$, and the second inequality must hold if $v_{i'}(\theta_{i'}) = v_i(\theta_i) - \epsilon$ for any $\epsilon > 0$. By substituting $v_{i'}(\theta_{i'})$ accordingly in the two inequalities and considering arbitrarily small $\epsilon > 0$ we obtain

$$s_i(j,\theta_i) = (\alpha_j - \alpha_{j+1}) \cdot v_i(\theta_i) + s_i(j+1,\theta_i).$$

Since this equality must hold for every agent, the payment for slot i equals

$$p_i = \sum_{j=i}^{\min(k,n-1)} (s_{j+1}(j,\theta_{j+1}) - s_{j+1}(j+1,\theta_{j+1})) = \sum_{j=i}^{\min(k,n-1)} ((\alpha_j - \alpha_{j+1}) \cdot v_{j+1}(\theta_{j+1})),$$

which is exactly the VCG payment for that slot.

The stronger claim for $n > k$ follows by setting $\alpha_{k+1} = 0$ and $s_i(k+1,\theta_i) = 0$, and deriving strategies inductively from slot k through 1 according to the above equality. □

In considering simplifications for combinatorial auctions, we have adopted the standard approach to assume incomplete information amongst agents, and in particular discussed a family of simplifications of the VCG mechanism that offers a tradeoff between social welfare and the amount of information agents have to communicate.

One may wonder why this tradeoff is necessary, and in how far it depends on the amount of information agents have about each others' types. It turns out that in the complete information case a much smaller number of bids is enough to preserve an efficient equilibrium, and in fact *all* equilibria, of the fully expressive VCG mechanism.

We show this using a simplification of the VCG mechanisms that we call n-VCG. The message space of n-VCG consists of all bid vectors with at most n non-zero entries, where n is the number of agents. This reduces the number of bids elicited from each agent from 2^k to at most n, which can be exponentially smaller. Surprisingly, this simplification is both tight and total, i.e., the set of Nash equilibria is completely unaffected by this restriction of the message space.

Theorem 8. *The n-VCG mechanism is tight and total with respect to Nash equilibria.*

Proof. We first prove that n-VCG is tight. By Proposition 1 it suffices to show that n-VCG satisfies outcome closure. Fix valuations v_j and types θ_j for all $j \in N$, and consider an arbitrary agent $i \in N$. We claim that for every $x_i \in X_i$ and every $\hat{x}_{-i} \in \hat{X}_{-i}$, there exists $\hat{x}_i \in \hat{X}_i$ such that

$$u_i((\hat{x}_i, \hat{x}_{-i}), \theta_i) \geq u_i((x_i, \hat{x}_{-i}), \theta_i).$$

Denote the outcome of VCG for $(x_i, \hat{x}_{-i})$ by o. Let $\hat{x}_i(C) = x_i(C)$ for $C = o_i$ and let $\hat{x}_i(C) = 0$ otherwise. We know that o achieves the same social welfare under $(\hat{x}_i, \hat{x}_{-i})$ as under $(x_i, \hat{x}_{-i})$. Since $\hat{x}_i(C) \leq x_i(C)$ for all C, we also know that the social welfare achieved by any outcome $o' \neq o$ under $(\hat{x}_i, \hat{x}_{-i})$ is weakly smaller than the social welfare achieved by o under $(\hat{x}_i, \hat{x}_{-i})$. Hence agent i gets the same bundle of items, namely o_i, under $(\hat{x}_i, \hat{x}_{-i})$ and $(x_i, \hat{x}_{-i})$. Since i's payment depends only on $\hat{x}_{-i}$ we conclude that $u_i((\hat{x}_i, \hat{x}_{-i}), \theta_i) \geq u_i((x_i, \hat{x}_{-i}), \theta_i)$.

Now we turn to totality. By Proposition 2 it suffices to show that n-VCG satisfies outcome reducibility. We compute $h(x)$ as follows: 1. Denote the outcome of VCG for x by o. For all agents i and $B = o_i$ mark $x_i(B)$. 2. For each agent j let o' denote the outcome of VCG if j was removed. For each such outcome o', all agents i, and $B = o'_i$ mark $x_i(B)$. 3. For all agents i and all bundles B set $\hat{x}_i(B) = x_i(B)$ if $x_i(B)$ was marked and $\hat{x}_i(B) = 0$ otherwise.

To (i): The outcome that maximizes welfare under x also maximizes welfare under $h(x)$ and the welfare achieved by this outcome is the same under x and $h(x)$. Similarly, for each agent i, the welfare achieved by the outcome that maximizes welfare if agent i is removed is the same under x and $h(x)$. Hence the outcomes $f(h(x))$ and $f(x)$ and the prices $p(h(x))$ and $p(x)$ for $h(x)$ and x are the same.

To (ii): Consider $\hat{x}'_i$. Denote the outcome computed by VCG for $(\hat{x}'_i, h_{-i}(x))$ by $\hat{o}$. We claim that there exists x'_i such that in the outcome o computed by VCG for (x'_i, x_{-i}) we have $o_i = \hat{o}_i$. Let $x'_i(C) = W_{max}(\hat{x}'_i, h_{-i}(x)) + \epsilon$ for $C = o_i$ and some $\epsilon > 0$ and let $x'_i(C) = 0$ otherwise. If $o_i = \hat{o}_i$, then the social welfare $W_{max}(x'_i, x_{-i})$ of the outcome of VCG for (x'_i, x_{-i}) is at least $W_{max}(\hat{x}'_i, h_{-i}(x)) + \epsilon$. Otherwise, the social welfare is $W_{max}(x'_i, x_{-i}) = W_{max}(0, x_{-i}) = W_{max}(0, h_{-i}(x)) \leq W_{max}(\hat{x}'_i, h_{-i}(x))$ and, thus, strictly smaller. We conclude that $o_i = \hat{o}_i$.

We further claim that the price $p_i(o_i)$ that agent i has to pay for o_i under (x'_i, x_{-i}) is at most the price $\hat{p}_i(o_i)$ that agent i has to pay for o_i under $(\hat{x}'_i, h_{-i}(x))$. Since $W_{max}(0, x_{-i}) = W_{max}(0, h_{-i}(x))$ and $W^{o_i}_{max}(0, x_{-i}) \geq W^{o_i}_{max}(0, h_{-i}(x))$, where $W^{o_i}_{max}(\ldots)$ denotes the maximal social welfare for the corresponding message profile if all items in o_i are removed from the set of items, we have that $p_i(o_i) = W_{max}(0, x_{-i}) - W^{o_i}_{max}(0, x_{-i}) \leq W_{max}(0, h_{-i}(x)) - W^{o_i}_{max}(0, h_{-i}(x)) = \hat{p}_i(o_i)$.

We conclude that $u_i((x'_i, x_{-i}), \theta_i) \geq u_i((\hat{x}'_i, h_{-i}(x)), \theta_i)$. □

2.7 Conclusion and Future Work

In this chapter we have studied simplifications of mechanisms obtained by restricting their message space, and have found that they can be used to solve different kinds of equilibrium selection problems that occur in practice. Direct revelation mechanisms typically have several equilibria, which might be more or less desirable from the point of view of the designer. Computational constraints might also imply that only a subset of the equilibria of a mechanisms can actually be achieved, in which case agents might have to select among several Pareto optimal equilibria. On the other hand, restricting the message space of a mechanism often reduces the amount of social welfare that can be achieved theoretically, and this seems to pivot on whether or not agents have complete information about each others' types. The choice between mechanisms with different degrees of expressiveness therefore involves a tradeoff between a benefit of simplicity and a price of simplicity.

The *price* of simplicity can easily be quantified, for example, by the loss in social welfare potentially incurred by a simplification, and has been studied in the context of both sponsored search and combinatorial auctions. Abrams et al. [1] and Blumrosen et al. [7], among others, have given bounds on the loss of social welfare incurred by α-GSP for different classes of valuations that are not proportional to the click-through rates. Christodoulou et al. [8], Bhawalkar and Roughgarden [5], and Feldman et al. [13] have studied the potential loss in welfare when a set of items is sold through a simplification of the combinatorial auction in which only bids that are additive in items are allowed. This bidding language requires only a small number of bids, and welfare in equilibrium turns out to be smaller by at most a constant factor in the number of agents than the optimum achieved by the fully expressive mechanism.

The *benefit* of simplicity is much harder to grasp as a concept. In this chapter we have argued that simplification can improve the economic properties of a mechanism by precluding bad or promoting good equilibria. A remaining challenge is to understand the benefit of simplicity in the context of simplified mechanisms for which the computation of an equilibrium might be an intractable problem, like the ones of Christodoulou et al., Bhawalkar and Roughgarden, and Feldman et al. described above. In contrast to the simplifications considered in the present paper, these mechanisms may not be able to solve the computational or informational problem of enabling agents to bid in a straightforward way. More generally, it is far from obvious how "straightforwardness" of a mechanism should be measured, but it seems reasonable to require that agents' strategies can be computed in polynomial time.

Bibliography

1. Z. Abrams, A. Ghosh, and E. Vee. Cost of conciseness in sponsored search auctions. In *Proceedings of the 3rd International Workshop on Internet and Network Economics*, volume

4858 of *Lecture Notes in Computer Science*, pages 326–334, 2007.

2. L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and P. Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, 2006.
3. M. Babaioff and T. Roughgarden. Equilibrium efficiency and price complexity in sponsored search auctions. In *Proceedings of the 6th Workshop on Ad Auctions*, 2010.
4. M. Benisch, N. Sadeh, and T. Sandholm. A theory of expressiveness in mechanisms. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, pages 17–23, 2008.
5. K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–709, 2011.
6. L. Blumrosen, N. Nisan, and I. Segal. Auctions with severely bounded communication. *Journal of Artificial Intelligence Research*, 28:233–266, 2007.
7. L. Blumrosen, J. Hartline, and S. Nong. Position auctions and non-uniform conversion rates. In *Proceedings of the 4th Workshop on Ad Auctions*, 2008.
8. G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, pages 820–832, 2008.
9. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
10. V. Conitzer and T. Sandholm. Computational criticisms of the revelation principle. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 262–263, 2004.
11. B. Edelman, M. Ostrovsky, and M. Schwartz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
12. M. Feldman and L. Blumrosen. Implementation with a bounded action space. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 62–71, 2006.
13. M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, page to appear, 2013.
14. R. D. Gomes and K. S. Sweeney. Bayes-Nash equilibria of the generalized second price auction. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 107–108, 2009.
15. T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

16. R. Holzman and D. Monderer. Characterization of ex post equilibrium in the VCG combinatorial auctions. *Games and Economic Behavior*, 47:87–103, 2004.

17. R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.

18. N. Hyafil and C. Boutilier. Partial revelation automated mechanism design. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 72–78, 2007.

19. N. Hyafil and C. Boutilier. Mechanism design with partial revelation. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1333–1340, 2007.

20. S. Lahaie and D. C. Parkes. On the communication requirements of verifying the VCG outcome. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 108–113, 2008.

21. S. Lahaie, D. M. Pennock, A. Saberi, and R. V. Vohra. Sponsored search auctions. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 28, pages 699–716. Cambridge University Press, 2007.

22. H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479, 1983.

23. P. Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70(1):62–70, 2010.

24. N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007.

25. N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006.

26. S. Shakkottai, R. Srikant, A. Ozdaglar, and D. Acemoglu. The price of simplicity. *IEEE Journal on Selected Areas in Communications*, 26(7):1269–1276, 2008.

27. H. Varian. Position auctions. *International Journal of Industrial Organization*, 25:1163–1178, 2007.

28. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

3 Payment Rules through Discriminant-Based Classifiers

3.1 Introduction

Mechanism design studies situations where a set of agents each hold private information about their preferences over different outcomes. The designer chooses a center that receives claims about such preferences, selects and enforces an outcome, and optionally collects payments. The classical approach is to impose *incentive compatibility*, ensuring that agents truthfully report their preferences in strategic equilibrium. Subject to this constraint, the goal is to identify a mechanism, i.e., a way of choosing an outcome and payments based on agents' reports, that optimizes a given design objective like social welfare, revenue, or some notion of fairness.

There are, however, significant challenges associated with this classical approach. First of all, it can be analytically cumbersome to derive optimal mechanisms for domains that are "multi-dimensional" in the sense that each agent's private information is described through more than a single number, and few results are known in this case. An example of a multi-dimensional domain is a combinatorial auction, where an agent's preferences are described by a value for each of several different bundles of items. Second, incentive compatibility can be costly, in that adopting it as a hard constraint can preclude mechanisms with useful economic properties. For example, imposing the strongest form of incentive compatibility, truthfulness in a dominant strategy equilibrium or *strategyproofness*, necessarily leads to poor revenue, vulnerability to collusion, and vulnerability to false-name bidding in combinatorial auctions where valuations exhibit complementarities among items [2, 23]. A third difficulty occurs when the optimal mechanism has an outcome or payment rule that is computationally intractable.

3.1.1 Our Contribution

In the face of these difficulties, we adopt statistical machine learning to automatically infer mechanisms with good incentive properties. Rather than imposing incentive compatibility as a hard constraint, we start from a given outcome rule, typically expressed as an algorithm, and then use machine learning techniques to identify a payment rule that minimizes agents' *expected ex post regret*. The ex post regret (or just *regret* where it causes no confusion) of an agent for truthful reporting in a given instance is the maximum amount by which its utility could be increased through a misreport holding constant the reports of others. The expected ex post regret is the average ex post regret over all agents and all preference types, calculated with respect to a distribution on types.

While a mechanism with zero regret for all agents on all inputs is strategyproof, we are especially interested in settings where the outcome rule does not allow for exact incentive compatibility. In this sense, the approach adopted here is not an equilibrium approach. But, there are two important comments to make in this regard. First, we insist that an agent's payment, conditioned on an outcome, is independent of its report. The only way an agent can improve its utility is by changing its report in a way that changes the outcome. Generically, this ensures mechanisms that provide zero marginal benefit to deviation from truthful reports. This property is seen in practice in the Generalized Second Price (GSP) mechanism used for sponsored search. This local stability property has been emphasized by Erdil and Klemperer [9] in the context of combinatorial auctions. In addition, a bound on expected regret implies a bound of the form "interim regret is at most ϵ with probability at least $1-\delta$," where interim regret is the ex post regret to an agent for a particular type, averaged over all types of other agents. Based on this, support for expected regret can be developed through a simple model of costly manipulation, where agents face some cost for trying to engage in strategic behavior, and choose not to engage in manipulation if this cost is greater than the expected gain. In this model, if the cost associated with strategic behavior is at least ϵ, an agent will find it beneficial to engage in manipulation with probability at most δ.

Our approach is applicable to domains that are multi-dimensional or for which the computational efficiency of outcome rules is a concern. Given the implied relaxation of incentive compatibility, the intended application is to domains in which incentive compatibility is unavailable or undesirable for outcome rules that meet certain economic and computational desiderata. The payment rule is learned on the basis of a given outcome rule, and as such the framework is most meaningful in domains where revenue considerations are secondary to outcome considerations.

The essential insight is that the payment rule of a strategyproof mechanism can be thought of as a classifier for predicting the outcome: the payment rule implies a price to each agent for each outcome, and the selected outcome must be one that simultaneously maximizes reported value minus price for every agent. The discriminant function of a classifier provides a score to different outcomes for a given input, with the outcome with the highest score corresponding to the prediction of the classier. By limiting classifiers to discriminant functions with this "value-minus-price" structure, where the price can be an arbitrary function of the outcome and the reports of other agents, we obtain a remarkably direct connection between multi-class classification and mechanism design. For an appropriate loss function, the discriminant function of a classifier that minimizes generalization error over a hypothesis class has a corresponding payment rule that minimizes expected ex post regret among all payment rules corresponding to classifiers in this class. Conveniently, an appropriate method exists for multi-class classification with large outcome spaces that supports the specific structure of the discriminant function, namely the method of *structural support vector machines* [26, 12]. Just like standard support vector machines, this also allows us to adopt non-linear kernels, thus enabling discriminant functions and thus price functions that depend in a non-linear way on the outcome and the reported types of agents.

In illustrating the framework, we focus on two situations where strategyproof payment rules are not available: a greedy outcome rule for a multi-minded combinatorial auction in which each agent is interested in a constant number of bundles, and an assignment problem with an egalitarian outcome rule, i.e., an outcome rule that maximizes the minimum value of any agent. The experimental results we obtain are encouraging, in that they demonstrate low expected ex post regret even when the 0/1 classification accuracy is only moderately good, and in particular better regret properties than those obtained through simple Vickrey-Clarke-Groves (VCG) based payment rules that we adopt as a baseline. In addition, we give special consideration to the failure of ex post individual rationality, and introduce methods to bias the classifier to avoid these kinds of errors as well as post hoc adjustments that eliminate them. As far as scalability is concerned, we emphasize that the computational cost associated with our approach occurs offline during training. The learned payment rules have a succinct description and can be evaluated quickly in a deployed mechanism.

3.1.2 Related Work

Conitzer and Sandholm [7] introduced the agenda of *automated mechanism design* (AMD) and formulated mechanism design as the search for an outcome rule and a payment rule among a class of rules satisfying incentive constraints. While the basic idea of optimal design is familiar from the seminal work of Myerson [20], a novel aspect of AMD is to formulate search over the space of all possible mappings from discrete type profiles to outcomes and payments. AMD is intractable when formulated on an explicit representation of the outcome and payment rules, because the type space is exponential in the number of agents.

One way to make AMD more tractable is to search through a parameterized space of incentive-compatible mechanisms [10]. More recently, advances in AMD have been made by considering domains with additive valuations and symmetry among agents, and by adopting Bayes-Nash incentive compatibility (BIC) rather than strategyproofness [5]. Still, these approaches seem limited to domains in which the outcome rule can be succinctly represented, which likely is not the case for the kinds of combinatorial auction problems considered here.

Lavi and Swamy [15] describe a method that takes any approximation algorithm for a set packing problem with a matching integrality gap and turns it into a mechanism with the same approximation guarantee that is strategyproof in expectation. Set packing includes combinatorial auctions as a special case. Bei and Huang [3] and Hartline et al. [11] describe an approach for turning an outcome rule into a mechanism that yields essentially the same expected amount of social welfare or social surplus and satisfies BIC. The approach computes an outcome and prices based on types sampled from probability distributions derived from the revealed types, and is applicable to both single-parameter and multi-parameter domains.

The target of minimizing expected ex post regret and the imposition of agent-independent prices make the incentive properties of mechanisms designed through our approach incomparable to BIC. On one hand, we are interested in minimizing statistics of *ex post* regret, and thus provide stronger guarantees than those of BIC. On the other hand, we don't guarantee zero expected regret (which would correspond to strategyproofness, and thus also BIC). Another distinction is that our approach can accommodate objectives that are non-separable across agents, such as in the egalitarian assignment problem.

In addition, in determining the outcome and payments for a given instance, the approach of Bei and Huang and Hartline et al. evaluates the outcome rule on a number of randomly perturbed replicas of that instance that is polynomial in the number of agents, the desired approximation ratio, and a notion capturing the complexity of the type spaces. When type spaces are large, as in the case of combinatorial auctions, this may become intractable. By contrast, our approach evaluates the outcome rule and the trained payment rule once for a given instance and incurs additional computational costs only during training.

The work of Lahaie [13, 14] precedes our work in adopting a kernel-based approach for combinatorial auctions, but focuses not on learning a payment rule for a given outcome rule but rather on solving the winner determination and pricing problem for a given instance of a combinatorial auction. Lahaie introduces the use of kernel methods to compactly represent non-linear price functions, which is also present in our work, but obtains incentive properties more indirectly through a connection between regularization and price sensitivity. The main difference is that Lahaie focuses on the design of scalable methods for clearing and pricing approximately welfare-maximizing combinatorial auctions, while we advance a framework for the automated design of payment rules that provide good incentive properties for a given outcome rule, which need not be welfare-maximizing.

Carroll [6] and Lubin and Parkes [19] provide surveys of related work on approximate incentive compatibility, or incentive compatibility in the large-market limit. A fair amount of attention has been devoted to regret-based metrics for quantifying the incentive properties of mechanisms (e.g., [21, 8, 17, 6]). Pathak and Sönmez [22] provide a qualitative ranking of different mechanisms without payments in terms of the number of manipulable instances. Budish [4] introduces an asymptotic, absolute design criterion regarding incentive properties in a large replica economy limit. Lubin and Parkes [18] provide experimental support that relates the divergence between the payoffs in a mechanism and the payoffs in a strategyproof "reference" mechanism to the amount by which agents deviate from truthful bidding in the Bayes-Nash equilibrium of a mechanism.

3.2 Preliminaries

A mechanism design problem is given by a set $N = \{1,2,\ldots,n\}$ of *agents* that interact to select an element from a set $\Omega \subseteq \times_{i\in N}\Omega_i$ of *outcomes*, where Ω_i denotes the set of possible outcomes for agent $i \in N$. Agent $i \in N$ is associated with a *type* θ_i from a set Θ_i of possible types, corresponding to the private information available to this agent. We write $\theta = (\theta_1,\ldots,\theta_n)$ for a profile of types for the different agents, $\Theta = \times_{i\in N}\Theta_i$ for the set of possible type profiles, and $\theta_{-i} \in \Theta_{-i}$ for a profile of types for all agents but i. Each agent $i \in N$ is further assumed to employ preferences over Ω_i, represented by a *valuation function* $v_i : \Theta_i \times \Omega_i \to \mathbb{R}$. We assume that for all $i \in N$ and $\theta_i \in \Theta_i$ there exists an outcome $o \in \Omega$ with $v_i(\theta_i, o_i) = 0$.

A *(direct) mechanism* is a pair (g, p) of an *outcome rule* $g : \Theta \to \times_{i\in N}\Omega_i$ and a *payment rule* $p : \Theta \to \mathbb{R}^n_{\geq 0}$. The intuition is that the agents reveal to the mechanism a type profile $\theta \in \Theta$, possibly different from their true types, and the mechanism chooses outcome $g(\theta)$ and charges each agent i a payment of $p_i(\theta) = (p(\theta))_i$. We assume *quasi-linear preferences*, so the *utility* of agent i with type $\theta_i \in \Theta_i$ given a profile $\theta' \in \Theta$ of revealed types is $u_i(\theta',\theta_i) = v_i(\theta_i, g_i(\theta')) - p_i(\theta')$, where $g_i(\theta) = (g(\theta))_i$ denotes the outcome for agent i. A crucial property of mechanism (g, p) is that its outcome rule is *feasible*, i.e., that $g(\theta) \in \Omega$ for all $\theta \in \Theta$.

Outcome rule g satisfies *consumer sovereignty* if for all $i \in N$, $o_i \in \Omega_i$, and $\theta'_{-i} \in \Theta_{-i}$, there exists $\theta'_i \in \Theta_i$ such that $g_i(\theta'_i,\theta'_{-i}) = o_i$; and *reachability of the null outcome* if for all $i \in N$, $\theta_i \in \Theta_i$, and $\theta'_{-i} \in \Theta_{-i}$, there exists $\theta'_i \in \Theta_i$ such that $v_i(\theta_i, g_i(\theta'_i,\theta'_{-i})) = 0$.

Mechanism (g, p) is *dominant strategy incentive compatible*, or *strategyproof*, if each agent maximizes its utility by reporting its true type, irrespective of the reports of the other agents, i.e., if for all $i \in N$, $\theta_i \in \Theta_i$, and $\theta' = (\theta'_i,\theta'_{-i}) \in \Theta$, $u_i((\theta_i,\theta'_{-i}),\theta_i) \geq u_i((\theta'_i,\theta'_{-i}),\theta_i)$; it satisfies *individual rationality* (IR) if agents reporting their true types are guaranteed non-negative utility, i.e., if for all $i \in N$, $\theta_i \in \Theta_i$, and $\theta'_{-i} \in \Theta_{-i}$, $u_i((\theta_i,\theta'_{-i}),\theta_i) \geq 0$. Observe that given reachability of the null outcome, strategyproofness implies individual rationality.

It is known that a mechanism (g, p) is strategyproof if and only if the payment of an agent is independent of its reported type and the chosen outcome simultaneously maximizes the utility of all agents, i.e., if for every $\theta \in \Theta$,

$$p_i(\theta) = t_i(\theta_{-i}, g_i(\theta)) \qquad \text{for all } i \in N, \text{ and} \tag{3.1}$$

$$g_i(\theta) \in \arg\max_{o'_i \in \Omega_i} \big(v_i(\theta_i, o'_i) - t_i(\theta_{-i}, o'_i)\big) \qquad \text{for all } i \in N, \tag{3.2}$$

for a *price function* $t_i : \Theta_{-i} \times \Omega_i \to \mathbb{R}$. This simple characterization is crucial for the main results in this chapter, providing the basis with which the discriminant function of a classifier can be used to induce a payment rule.

In addition, a direct characterization of strategyproofness in terms of monotonicity properties of outcome rules explains which outcome rules can be associated with a payment rule in order to be "implementable" within a strategyproof mechanism [24, 1]. These monotonicity properties provide a fundamental constraint on when our machine learning framework can hope to identify a payment rule that provides full strategyproofness.

We quantify the degree of strategyproofness of a mechanism in terms of the *regret* experienced by an agent when revealing its true type, i.e., the potential gain in utility by revealing a different type instead. Formally, the *ex post regret* of agent $i \in N$ in mechanism (g, p), given true type $\theta_i \in \Theta_i$ and reported types $\theta'_{-i} \in \Theta_{-i}$ of the other agents, is

$$rgt_i(\theta_i, \theta'_{-i}) = \max_{\theta'_i \in \Theta_i} u_i\big((\theta'_i, \theta'_{-i}), \theta_i\big) - u_i\big((\theta_i, \theta'_{-i}), \theta_i\big).$$

Analogously, the *ex post violation of individual rationality* of agent $i \in N$ in mechanism (g, p), given true type $\theta_i \in \Theta_i$ and reported types $\theta'_{-i} \in \Theta_{-i}$ of the other agents, is

$$irv_i(\theta_i, \theta'_{-i}) = |\min(u_i((\theta_i, \theta'_{-i}), \theta_i), 0)|.$$

We consider situations where types are drawn from a distribution with probability density function $D : \Theta \to \mathbb{R}$ such that $D(\theta) \geq 0$ and $\int_{\theta \in \Theta} D(\theta) = 1$. Given such a distribution, and assuming that all agents report their true types, the *expected ex post regret* of agent $i \in N$ in mechanism (g, p) is $\mathbb{E}_{\theta \sim D}[rgt_i(\theta_i, \theta_{-i})]$.

Outcome rule g is *agent symmetric* if for every permutation π of N and all types $\theta, \theta' \in \Theta$ such that $\theta_i = \theta'_{\pi(i)}$ for all $i \in N$, $g_i(\theta) = g_{\pi(i)}(\theta')$ for all $i \in N$. Note that this specifically requires that $\Theta_i = \Theta_j$ and $\Omega_i = \Omega_j$ for all $i, j \in N$. Similarly, type distribution D is *agent symmetric* if $D(\theta) = D(\theta')$ for every permutation π of N and all types $\theta, \theta' \in \Theta$ such that $\theta_i = \theta'_{\pi(i)}$ for all $i \in N$. Given agent symmetry, a price function $t_1 : \Theta_{-1} \times \Omega_i \to \mathbb{R}$ for agent 1 can be used to generate the payment rule p for a mechanism (g, p), with

$$p(\theta) = \big(t_1(\theta_{-1}, g_1(\theta)), t_1(\theta_{-2}, g_2(\theta)), \ldots, t_1(\theta_{-n}, g_n(\theta))\big),$$

so that the expected ex post regret is the same for every agent.

We assume agent symmetry in the sequel, which precludes outcome rules that break ties based on agent identity, but obviates the need to train a separate classifier for each agent while also providing some benefits in terms of presentation. Because ties occur only with negligible probability in our experimental framework, the experimental results are not affected by this assumption.

3.3 Payment Rules from Multi-Class Classifiers

A *multi-class classifier* is a function $h : X \to Y$, where X is an input domain and Y is a discrete output domain. One could imagine, for example, a multi-class classifier that labels a given image as that of a dog, a cat, or some other animal. In the context of mechanism design, we will be interested in classifiers that take as input a type profile and output an outcome. What distinguishes this from an outcome rule is that we will impose restrictions on the form the classifier can take.

Classification typically assumes an underlying target function $h^* : X \to Y$, and the goal is to learn a classifier h that minimizes disagreements with h^* on a given input distribution D on X, based only on a finite set of *training data* $\{(x^1, y^1), \ldots, (x^\ell, y^\ell)\} = \{(x^1, h^*(x^1)), \ldots, (x^\ell, h^*(x^\ell))\}$ with $x^1, \ldots, x^\ell$ drawn from D. If $h(x) = h^*(x)$ for all $x \in X$, we say that h is a *perfect classifier* for h^*.

The classifier h typically comes from a certain class of functions, referred to as the *hypothesis class* $\mathscr{H}$. We consider classifiers that are defined in terms of a *discriminant function* $f : X \times Y \to \mathbb{R}$, which assigns a score to each input-output pair. The corresponding discriminant-based classifier h chooses for a given input x an output y with maximal score, i.e.,

$$h(x) \in \arg\max_{y \in Y} f(x, y)$$

for all $x \in X$. More specifically, we will be concerned with *linear* discriminant functions of the form

$$f_w(x, y) = w^T \psi(x, y)$$

for a weight vector $w \in \mathbb{R}^m$ and a *feature map* $\psi : X \times Y \to \mathbb{R}^m$, where $m \in \mathbb{N} \cup \{\infty\}$.[1] The function ψ maps input-output pairs into an m-dimensional space in a not necessarily linear manner. Hence, although f_w is linear in $\psi(x, y)$, it need not be linear in (x, y), which generally allows for non-linear classification.

3.3.1 Mechanism Design as Classification

Assume that we are given an outcome rule g and access to a distribution D over type profiles, and want to design a corresponding payment rule p that gives the mechanism (g, p) the best possible incentive properties. Assuming agent symmetry, we focus on a partial outcome rule $g_1 : \Theta \to \Omega_1$ and train a classifier to predict the outcome to agent 1. To train a classifier, we generate examples by drawing a type profile $\theta \in \Theta$ from distribution D and applying outcome rule g to obtain the target class $g_1(\theta) \in \Omega_1$.

We impose a special structure on the hypothesis class. A classifier $h_w : \Theta \to \Omega_1$ is *admissible* if it is defined in terms of a discriminant function f_w of the form

$$f_w(\theta, o_1) = w_1 v_1(\theta_1, o_1) + w_{-1}^T \psi(\theta_{-1}, o_1)$$

for weights w such that $w_1 \in \mathbb{R}_{>0}$ and $w_{-1} \in \mathbb{R}^m$, and a feature map $\psi : \Theta_{-1} \times \Omega_1 \to \mathbb{R}^m$ for

[1] We allow w to have infinite dimension, but require the inner product between w and $\psi(x, y)$ to be defined in any case. Computationally the infinite-dimensional case is handled through the kernel trick, which is described in Section 3.4.1.

$m \in \mathbb{N} \cup \{\infty\}$.

Only the first term of $f_w(\theta, o_1)$ depends on the type θ_1 of agent 1, while the remaining terms ignore agent 1's type θ_1 entirely. This restriction allows us to directly infer agent-independent prices from a trained classifier. For this, define the *associated price function* of an admissible classifier h_w as

$$t_w(\theta_{-1}, o_1) = -\frac{1}{w_1} w_{-1}^T \psi(\theta_{-1}, o_1),$$

where we again focus on agent 1 for concreteness. By agent symmetry, we obtain the mechanism (g, p_w) corresponding to classifier h_w by letting

$$p_w(\theta) = \big(t_w(\theta_{-1}, g_1(\theta)), t_w(\theta_{-2}, g_2(\theta)), \ldots, t_w(\theta_{-n}, g_n(\theta))\big).$$

Even with admissibility, appropriate choices for the feature map ψ will produce rich families of classifiers, and thus ultimately useful payment rules. Moreover, this form is compatible with structural support vector machines, discussed in Section 3.4.1.

3.3.2 Example: Single-Item Auction

Before proceeding further, we illustrate the ideas developed so far in the context of a single-item auction. In a single-item auction, the type of each agent is a single number, corresponding to its value for the item being auctioned, and there are two possible allocations from the point of view of agent 1: one where it receives the item, and one where it does not. Formally, $\Theta = \mathbb{R}^n$ and $\Omega_1 = \{0, 1\}$.

Consider a setting with three agents and a training set

$$(\theta^1, o_1^1) = ((1,3,5), 0), \quad (\theta^2, o_1^2) = ((5,4,3), 1), \quad (\theta^3, o_1^3) = ((2,3,4), 0),$$

and note that this training set is consistent with an *optimal* outcome rule, i.e., one that assigns the item to an agent with maximum value. Our goal is to learn an admissible classifier

$$h_w(\theta) = \arg\max_{o_1 \in \{0,1\}} f_w(\theta, o_1) = \arg\max_{o_1 \in \{0,1\}} w_1 v_1(\theta_1, o_1) + w_{-1}^T \psi(\theta_{-1}, o_1)$$

that performs well on the training set. Since there are only two possible outcomes, the outcome chosen by h_w is simply the one with the larger discriminant. A classifier that is perfect on the training data must therefore satisfy the following constraints:

$$\begin{aligned}
w_1 \cdot 0 + w_{-1}^T \psi((3,5), 0) &> w_1 \cdot 1 + w_{-1}^T \psi((3,5), 1), \\
w_1 \cdot 5 + w_{-1}^T \psi((4,3), 1) &> w_1 \cdot 0 + w_{-1}^T \psi((4,3), 0), \\
w_1 \cdot 0 + w_{-1}^T \psi((3,4), 0) &> w_1 \cdot 2 + w_{-1}^T \psi((3,4), 1).
\end{aligned}$$

This can for example be achieved by setting $w_1 = 1$ and

$$w_{-1}^T \psi((\theta_2,\theta_3),o_1) = \begin{cases} -\max(\theta_2,\theta_3) & \text{if } o_1 = 1 \text{ and} \\ 0 & \text{if } o_1 = 0. \end{cases} \tag{3.3}$$

Recalling our definition of the price function as $t_w(\theta_{-1},o_1) = -(1/w_1)\, w_{-1}^T \psi(\theta_{-1},o_1)$, we see that this choice of w and ψ corresponds to the second-price payment rule. We will see in the next section that this relationship is not a coincidence.[2]

3.3.3 Perfect Classifiers and Implementable Outcome Rules

We now formally establish a connection between implementable outcome rules and perfect classifiers.

Theorem 1. *Let (g,p) be a strategyproof mechanism with an agent symmetric outcome rule g, and let t_1 be the corresponding price function. Then, a perfect admissible classifier h_w for partial outcome rule g_1 exists if* $\arg\max_{o_1\in\Omega_1}(v_1(\theta_1,o_1) - t_1(\theta_{-1},o_1)))$ *is unique.*

Proof. By the first characterization of strategyproof mechanisms, g must select an outcome that maximizes the utility of agent 1 at the current prices, i.e.,

$$g_1(\theta) \in \arg\max_{o_1\in\Omega_1}(v_1(\theta_i,o_1) - t_1(\theta_{-1},o_1)).$$

Consider the admissible discriminant $f_{(1,1)}(\theta,o_1) = v_1(\theta_1,o_1) - t_1(\theta_{-1},o_1)$, which uses the price function t_1 as its feature map. Clearly, the corresponding classifier $h_{(1,1)}$ maximizes the same quantity as g_1, and the two must agree if there is a unique maximizer. □

The relationship also works in the opposite direction: a perfect, admissible classifier h_w for outcome rule g can be used to construct a payment rule that turns g into a strategyproof mechanism.

Theorem 2. *Let g be an agent symmetric outcome rule, $h_w : \Theta \to \Omega_1$ an admissible classifier, and p_w the payment rule corresponding to h_w. If h_w is a perfect classifier for the partial outcome rule g_1, then the mechanism (g,p_w) is strategyproof.*

We prove this result by expressing the regret of an agent in mechanism (g,p_w) in terms of the discriminant function f_w. Let $\Omega_i(\theta_{-i}) \subseteq \Omega_i$ denote the set of partial outcomes for agent i that can be obtained under g given reported types θ_{-i} from all agents but i, keeping the dependence on g silent for notational simplicity.

[2] In practice, we are limited in the machine learning framework to hypotheses that are linear in $\psi((\theta_2,\theta_3),o_1)$, and will not be able to guarantee that (3.3) holds exactly. In Section 3.4.1 we will see, however, that certain choices of ψ allow for very complex hypotheses that can closely approximate arbitrary functions.

Lemma 1. *Suppose that agent* 1 *has type* θ_1 *and that the other agents report types* θ_{-1}. *Then the regret of agent* 1 *for bidding truthfully in mechanism* (g, p_w) *is*

$$\frac{1}{w_1}\left(\max_{o_1\in\Omega(\theta_{-1})} f_w(\theta, o_1) - f_w(\theta, g_1(\theta))\right).$$

Proof. We have

$$\begin{aligned} rgt_1(\theta) = & \max_{\theta_1'\in\Theta_1} \big(v_1(\theta_1, g_1(\theta_1', \theta_{-1})) - p_{w,1}(\theta_1', \theta_{-1})\big) - \big(v_1(\theta_1, g_1(\theta)) - p_{w,1}(\theta)\big) \\ = & \max_{o_1\in\Omega_1(\theta_{-1})} \big(v_1(\theta_1, o_1) - t_w(\theta_{-1}, o_1)\big) - \big(v_1(\theta_1, g_1(\theta)) - t_w(\theta_{-1}, g_1(\theta))\big) \\ = & \max_{o_1\in\Omega_1(\theta_{-1})} \big(v_1(\theta_1, o_1) + \frac{1}{w_1} w_{-1}^T \psi(\theta_{-1}, o_1)\big) - \big(v_1(\theta_1, g_1(\theta)) + \frac{1}{w_1} w_{-1}^T \psi(\theta_{-1}, g_1(\theta))\big) \\ = & \frac{1}{w_1}\big(\max_{o_1\in\Omega_1(\theta_{-1})} f_w(\theta, o_1) - f_w(\theta, g_1(\theta))\big). \end{aligned}$$

□

Proof of Theorem 2. If h_w is a perfect classifier, then the discriminant function f_w satisfies $\arg\max_{o_1\in\Omega_1} f_w(\theta, o_1) = g_1(\theta)$ for every $\theta \in \Theta$. Since $g_1(\theta) \in \Omega_1(\theta_{-1})$, we thus have that $\max_{o_1\in\Omega_1(\theta_{-1})} f_w(\theta, o_1) = f_w(\theta, g_1(\theta))$. By Lemma 1, the regret of agent 1 for bidding truthfully in mechanism (g, p_w) is always zero, which means that the mechanism is strategyproof. □

It bears emphasis that classifier h_w is only used to derive the payment rule p_w, while the outcome is still selected according to g. In principle, classifier h_w could be used to obtain an agent symmetric outcome rule g_w and, since h_w is a perfect classifier for itself, a strategyproof mechanism (g_w, p_w). Unfortunately, outcome rule g_w is not in general feasible. Mechanism (g, p_w), on the other hand, is not strategyproof when h_w fails to be a perfect classifier for g. While payment rule p_w always satisfies the agent-independence property (3.1) required for strategyproofness, the "optimization" property (3.2) might be violated when $h_w(\theta) \neq g_1(\theta)$.

3.3.4 Approximate Classification and Approximate Strategyproofness

A perfect admissible classifier for outcome rule g leads to a payment rule that turns g into a strategyproof mechanism. We now show that this result extends gracefully to situations where no such payment rule is available, by relating the *expected* ex post regret of a mechanism (g, p) to a measure of the generalization error of a classifier for g.

Fix a feature map ψ, and denote by $\mathcal{H}_\psi$ the space of all admissible classifiers with this feature map. The *discriminant loss* of a classifier $h_w \in \mathcal{H}_\psi$ with respect to a type profile θ and an outcome $o_1 \in \Omega_1$ is given by

$$\Delta_w(o_1, \theta) = \frac{1}{w_1}\big(f_w(\theta, h_w(\theta)) - f_w(\theta, o_1)\big).$$

Intuitively the discriminant loss measures how far, in terms of the normalized discriminant,

h_w is from predicting the correct outcome for type profile θ, assuming the correct outcome is o_1. Note that for all $\theta \in \Theta$ we have $\Delta(o_1, \theta) \geq 0$ for all $o_1 \in \Omega_1$, and $\Delta(o_1, \theta) = 0$ if $o_1 = h_w(\theta)$. Note further that for all $\theta \in \Theta$ and any two $h_w, h_{w'} \in \mathcal{H}_\psi$ the fact that $h_w(\theta) = h_{w'}(\theta)$ does not imply that $\Delta_w(o_1, \theta) = \Delta_{w'}(o_1, \theta)$ for all $o_1 \in \Omega_1$: even if two classifiers predict the same outcome, one of them may still be closer to predicting the correct outcome.

The *generalization error* of classifier $h_w \in \mathcal{H}_\psi$ with respect to a type distribution D and a partial outcome rule $g_1 : \Theta \to \Omega_1$ is then given by

$$R_w(D, g) = \int_{\theta \in \Theta} \Delta_w\big(g_1(\theta), \theta\big) D(\theta)\, d\theta.$$

The following result establishes a connection between the generalization error and the expected ex post regret of the corresponding mechanism.

Theorem 3. *Consider an outcome rule g, a space $\mathcal{H}_\psi$ of admissible classifiers, and a type distribution D. Let $h_{w^*} \in \mathcal{H}_\psi$ be a classifier that minimizes generalization error with respect to D and g among all classifiers in $\mathcal{H}_\psi$. Then the following holds:*

1. *If g satisfies consumer sovereignty, then (g, p_{w^*}) minimizes expected ex post regret with respect to D among all mechanisms (g, p_w) corresponding to classifiers $h_w \in \mathcal{H}_\psi$.*
2. *Otherwise, (g, p_{w^*}) minimizes an upper bound on expected ex post regret with respect to D amongst all mechanisms (g, p_w) corresponding to classifiers $h_w \in \mathcal{H}_\psi$.*

Proof. For the second property, observe that

$$\begin{aligned}
\Delta_w(g_1(\theta), \theta) &= \frac{1}{w_1}\big(f_w(\theta, h_w(\theta)) - f_w(\theta, g_1(\theta))\big) \\
&= \frac{1}{w_1}\big(\max_{o_1 \in \Omega_1} f_w(\theta, o_1) - f_w(\theta, g_1(\theta))\big) \\
&\geq \frac{1}{w_1}\big(\max_{o_1 \in \Omega(\theta_{-1})} f_w(\theta, o_1) - f_w(\theta, g_1(\theta))\big) \\
&= rgt_1(\theta),
\end{aligned}$$

where the last equality holds by Lemma 1. If g satisfies consumer sovereignty, then the inequality holds with equality, and the first property follows as well. □

It should be noted at this point that while exact classifiers h_w generally induce mechanisms (g, p_w) that satisfy individual rationality, this ceases to be true for mechanisms (g, p_w) induced by inexact classifiers h_w. We discuss ways to bias the classifier to avoid errors that lead to IR violations as well as post hoc adjustments that eliminate them in Section 3.4.2.

3.4 A Solution using Structural Support Vector Machines

In this section we discuss the method of *structural support vector machines* (structural SVMs) [26, 12], and show how it can be adapted for the purpose of learning classifiers with admissible discriminant functions.

3.4.1 Structural SVMs

Given an input space X, a discrete output space Y, a target function $h^* : X \to Y$, and a set of *training examples* $\{(x^1, h^*(x^1)), \ldots, (x^\ell, h^*(x^\ell))\} = \{(x^1, y^1), \ldots, (x^\ell, y^\ell)\}$, structural SVMs learn a multi-class classifier h that on input $x \in X$ selects an output $y \in Y$ that maximizes $f_w(x, y) = w^T \psi(x, y)$. For a given feature map ψ, the training problem is to find a vector w for which h_w has low generalization error.

Given examples $\{(x^1, y^1), \ldots, (x^\ell, y^\ell)\}$ and a parameter C, training is achieved by solving the following convex optimization problem:

$$\min_{w,\xi\geq 0} \frac{1}{2} w^T w + \frac{C}{\ell} \sum_{k=1}^{\ell} \xi^k \qquad \text{(Training Problem 1)}$$

$$\text{s.t. } w^T\big(\psi(x^k, y^k) - \psi(x^k, y)\big) \geq \mathscr{L}(y^k, y) - \xi^k \quad \text{for all } k = 1, \ldots, \ell,\ y \in Y$$

$$\xi^k \geq 0 \quad \text{for all } k = 1, \ldots, \ell.$$

The goal is to find a weight vector w and slack variables ξ^k such that the objective function is minimized while satisfying the constraints. The learned weight vector w parameterizes the discriminant function f_w, which in turn defines the classifier h_w. The k-th constraint states that the value of the discriminant function on (x^k, y^k) should exceed the value of the discriminant function on (x^k, y) by at least $\mathscr{L}(y^k, y)$ for all $y \in Y$, where $\mathscr{L}$ is a loss function that penalizes misclassification, with $\mathscr{L}(y, y) = 0$ and $\mathscr{L}(y, y') \geq 0$ for all $y, y' \in Y$. We generally use a 0/1 loss function, but consider an alternative in Section 3.4.2 to improve ex post IR properties. Positive values for the slack variables ξ^k allow the weight vector to violate some of the constraints.

The other term in the objective, the squared norm of w, penalizes scaling of w. This is necessary because scaling of w can arbitrarily increase the margin between $f_w(x^k, y^k)$ and $f_w(x^k, y)$ and make the constraints easier to satisfy. Smaller values of w, on the other hand, increases the ability of the learned classifier to generalize by decreasing the propensity to over-fit to the training data. Parameter C is therefore a regularization parameter: larger values of C encourage small ξ^k and larger w, such that more points are classified correctly, but with a smaller margin.

The Feature Map and the Kernel Trick

Given a feature map ψ, the *feature vector* $\psi(x, y)$ for $x \in X$ and $y \in Y$ provides an alternate representation of the input-output pair (x, y). It is useful to consider feature maps ψ for which $\psi(x, y) = \phi(\chi(x, y))$, where $\chi : X \times Y \to \mathbb{R}^s$ for some $s \in \mathbb{N}$ is an *attribute map* that combines x and y into a single *attribute vector* $\chi(x, y)$ compactly representing the pair, and $\phi : \mathbb{R}^s \to \mathbb{R}^m$ for $m > s$ maps the attribute vector to a higher-dimensional space in a non-linear way. In this way, SVMs can achieve non-linear classification in the original space.

While we work hard to keep s small, the so-called *kernel trick* means that we do not have the same problem with m: it turns out that in the dual of Training Problem 1, $\psi(x, y)$ only appears in an inner product of the form $\langle\psi(x, y), \psi(x', y')\rangle$, or, for a decomposable feature map, $\langle\phi(z), \phi(z')\rangle$ where $z = \chi(x, y)$ and $z' = \chi(x', y')$. For computational tractability it therefore suffices that this inner product can be computed efficiently, and the "trick" is to choose ϕ such that $\langle\phi(z), \phi(z')\rangle = K(z, z')$ for a simple closed-form function K, known as the *kernel*.

We consider *polynomial kernels* K_{polyd}, parameterized by $d \in \mathbb{N}^+$, and *radial basis function (RBF) kernels* K_{RBF}, parameterized by $\gamma = 1/(2\sigma^2)$ for $\sigma \in \mathbb{R}^+$:

$$K_{polyd}(z, z') = (z \cdot z')^d,$$
$$K_{RBF}(z, z') = \exp\left(-\gamma\left(\|z\|^2 + \|z'\|^2 - 2z \cdot z'\right)\right).$$

Both polynomial and RBF kernels use the standard inner product of their arguments, so their efficient computation requires that $\chi(x, y) \cdot \chi(x, y')$ can be computed efficiently.

Dealing with an Exponentially Large Output Space

Training Problem 1 has $\Omega(|Y|\ell)$ constraints, where Y is the output space and ℓ the number of training instances, and enumerating all of them is computationally prohibitive when Y is large. Joachims et al. [12] address this issue for structural SVMs through constraint generation: starting from an empty set of constraints, this technique iteratively adds a constraint that is maximally violated by the current solution until that violation is below a desired threshold ϵ. Joachims et al. show that this will happen after no more than $O(\frac{C}{\epsilon})$ iterations, each of which requires $O(\ell)$ time and memory. However, this approach assumes the existence of an efficient separation oracle, which given a weight vector w and an input x finds an output $y \in \arg\max_{y' \in Y} f_w(x, y')$. The existence of such an oracle remains an open question in application to combinatorial auctions; see Section 3.5.1 for additional discussion.

Required Information

In summary, the use of structural SVMs requires specification of the following:

1. The input space X, the discrete output space Y, and examples of input-output pairs.

2. An attribute map $\chi : X \times Y \to \mathbb{R}^s$.This function generates an attribute vector that combines the input and output data into a single object.
3. A kernel function $K(z, z')$, typically chosen from a well-known set of candidates, e.g., polynomial or RBF. The kernel implicitly calculates the inner product $\langle \phi(z), \phi(z') \rangle$, e.g., between a mapping of the inputs into a high dimensional space.
4. If the space Y is prohibitively large, a routine that allows for efficient separation, i.e., a function that computes $\arg\max_{y \in Y} f_w(x, y)$ for a given w, x.

In addition, the user needs to stipulate particular training parameters, such as the regularization parameter C, and the kernel parameter γ if the RBF kernel is being used.

3.4.2 Structural SVMs for Mechanism Design

We now specialize structural SVMs such that their learned discriminant function will manifest as a payment rule for a given symmetric outcome function g and distribution D. For this we specify the input domain X and the output domain Y, and we impose a special structure on the attribute map ψ.

The input domain X is the space of type profiles Θ, and the output domain Y is the space Ω_1 of outcomes for agent 1. Thus we construct training data by sampling $\theta \sim D$ and applying g to these inputs: $\{(\theta^1, g_1(\theta^1)), \ldots, (\theta^\ell, g_1(\theta^\ell))\} = \{(\theta^1, o_1^1), \ldots, (\theta^\ell, o_1^\ell)\}$. For admissibility of the learned hypothesis $h_w(\theta) = \arg\max_{o_1 \in \Omega_1} w^T \psi(\theta, o_1)$, we require that

$$\psi(\theta, o_1) = (v_1(\theta_1, o_1), \psi'(\theta_{-1}, o_1))$$

When learning payment rules, we therefore use an attribute map $\chi' : \Theta_{-1} \times \Omega_1 \to \mathbb{R}^s$ rather than $\chi : \Theta \times \Omega_1 \to \mathbb{R}^s$, and the kernel ϕ' we specify will only be applied to the output of χ'.

This results in the following more specialized training problem:

$$\min_{w, \xi \geq 0} \frac{1}{2} w^T w + \frac{C}{\ell} \sum_{k=1}^{\ell} \xi^k \qquad \text{(Training Problem 2)}$$

$$\text{s.t. } (w_1 v_1(\theta_1^k, o_1^k) + w_{-1}^T \psi'(\theta_{-1}^k, o_1^k)) - (w_1 v_1(\theta_1^k, o_1) + w_{-1}^T \psi'(\theta_{-1}^k, o_1)) \geq \mathcal{L}(o_1^k, o_1) - \xi^k$$
$$\text{for all } k = 1, \ldots, \ell,\ o_1 \in \Omega_1$$

$$\xi^k \geq 0 \quad \text{for all } k = 1, \ldots, \ell.$$

If $w_1 > 0$ then the weights w together with the feature map ψ' define a price function $t_w(\theta_{-1}, o_1) = -(1/w_1) w_{-1}^T \psi'(\theta_{-1}, o_1)$ that can be used to define payments $p_w(\theta)$, as described in Section 3.3.1. In this case, we can also relate the regret in the induced mechanism (g, p_w) to the classification error as described in Section 3.3.3.

Theorem 4. *Consider training data* $\{(\theta^1, o_1^1), \ldots, (\theta^\ell, o_1^\ell)\}$. *Let* g *be an outcome function such that* $g_1(\theta^k) = o_1^k$ *for all* k. *Let* w, ξ^k *be the weight vector and slack variables output by Training*

Problem 2, with $w_1 > 0$*. Consider the corresponding mechanism* (g, p_w)*. For each* θ^k*,*

$$rgt_1(\theta^k) \leq \frac{1}{w_1}\xi^k$$

Proof. Consider input θ^k. The constraints in the training problem impose that for every outcome $o_1 \in \Omega_1$,

$$w_1 v_1(\theta_1^k, o_1^k) + w_{-1}^T \psi'(\theta_{-1}^k, o_1^k) - \left(w_1 v_1(\theta_1^k, o_1) + w_{-1}^T \psi'(\theta_{-1}^k, o_1)\right) \geq \mathscr{L}(o_1^k, o_1) - \xi^k$$

Rearranging,

$$\begin{aligned} &\xi^k \geq \mathscr{L}(o_1^k, o_1) + \left(w_1 v_1(\theta_1^k, o_1) + w_{-1}^T \psi'(\theta_{-1}^k, o_1)\right) - \left(w_1 v_1(\theta_1^k, o_1^k) + w_{-1}^T \psi'(\theta_{-1}^k, o_1^k)\right) \\ \Rightarrow &\xi^k \geq \mathscr{L}(o_1^k, o_1) + f_w(\theta^k, o_1) - f_w(\theta^k, o_1^k) \end{aligned}$$

This inequality holds for every $o_1 \in \Omega_1$, so

$$\begin{aligned} \xi^k &\geq \max_{o_1 \in \Omega_1} \left(\mathscr{L}(o_1^k, o_1) + f_w(\theta^k, o_1) - f_w(\theta^k, o_1^k)\right) \\ &\geq \max_{o_1 \in \Omega_1} \left(f_w(\theta^k, o_1) - f_w(\theta^k, o_1^k)\right) \\ &\geq w_1 \, rgt_1(\theta^k) \end{aligned}$$

where the second inequality holds because $\mathscr{L}(o_1^k, o_1) \geq 0$, and the final inequality follows from Lemma 1. This completes the proof. □

We choose not to enforce $w_1 > 0$ explicitly in Training Problem 2, because adding this constraint leads to a dual problem that references ψ' outside of an inner product and thus makes computation of all but linear or low-dimensional polynomial kernels prohibitively expensive. Instead, in our experiments we simply discard hypotheses where the result of training is $w_1 \leq 0$. This is sensible since the discriminant function value should increase as an agent's value increases, and negative values of w_1 typically mean that the training parameter C or the kernel parameter γ (if the RBF kernel is used) are poorly chosen. It turns out that w_1 is indeed positive most of the time, and for every experiment a majority of the choices of C and γ yield positive w_1 values. For this reason, we do not expect the requirement that $w_1 > 0$ to be a problem in practice.[3]

Payment Normalization

One issue with the framework as stated is that the payments p_w computed from the solution to Training Problem 2 could be negative.

[3] For multi-minded combinatorial auctions, $1049/1080 > 97\%$ of the trials had positive w_1, for the assignment problem all of the trials did; see Section 3.5 for details.

We solved this problem by normalizing payments, using a *baseline outcome* o_b: if there exists an outcome o' such that $v_1(\theta_1, o') = 0$ for every θ_1, this "null outcome" is used as the baseline; otherwise, we use the outcome with the lowest payment. Let $t_w(\theta_{-1}, o_1)$ be the price function corresponding to the solution w to Training Problem 2. Adopting the baseline outcome, the *normalized payments* $t'_w(\theta_{-1}, o_1)$ are defined as

$$t'_w(\theta_{-1}, o_1) = \max(0, t_w(\theta_{-1}, o_1) - t_w(\theta_{-1}, o_b)).$$

Note that o_b is only a function of θ_{-1}, even when there is no null outcome, so t'_w is still only a function of θ_{-1} and o_1.

Individual Rationality Violation

Even after normalization, the learned payment rule p_w may not satisfy IR. We offer three solutions to this problem, which can be used in combination.

Payment offsets One way to decrease the rate of IR violation is to add a *payment offset*, which decreases all payments (for all type reports) by a given amount. We apply this payment offset to all bundles other than o_b; as with payment normalization, the adjusted payment is set to 0 if it is negative.[4] Note that payment offsets decrease IR violation, but may increase regret. For instance, suppose there are only two outcomes o_{11}, o_{12}, where o_{12} is the null outcome. Suppose agent 1 values o_{11} at 5 and receives the null outcome if he reports truthfully. Suppose further that payments t_w are 7 for o_{11} and 0 for the null outcome. With no payment offset, the agent experiences no regret, since he receives utility 0 from the null outcome, but negative utility from o_{11}. However, if the payment offset is greater than 2, the agent's regret becomes positive (assuming consumer sovereignty) because he could have reported differently and received o_{11} and received positive utility.

Adjusting the loss function We incur an IR violation when there is a null outcome o_{null} such that $g_1(\theta) \neq o_{null}$ and $f_w(\theta, o_{null}) > f_w(\theta, g_1(\theta))$ for some type θ, assuming truthful reports. This happens because $f_w(\theta, o_1)$ is a scaled version of the agent's utility for outcome o_1 under payments p_w. If the utility for the null outcome is greater than the utility for $g_1(\theta)$, then the payment $t_w(\theta_{-1}, g_1(\theta))$ must be greater than $v_1(\theta_1, g_1(\theta))$, causing an IR violation. We can discourage these types of errors by modifying the constraints of Training Problem 2: when $o_1^k \neq o_{null}$ and $o_1 = o_{null}$, we can increase $\mathscr{L}(o_1^k, o_1)$ to heavily penalize misclassifications of this type. With a larger $\mathscr{L}(o_1^k, o_1)$, a larger ξ^k will be required if $f_w(\theta, o_1^k) < f_w(\theta, o_{null})$. As with payment offsets, this technique will decrease IR violations but is not guaranteed to eliminate all of them. In our experimental results, we refer to this as the *null loss fix*, and the null loss refers to the value we choose for $\mathscr{L}(o_1^k, o_{null})$ where $o_1^k \neq o_{null}$.

[4] It is again crucial that o_b depends only on θ_{-1}, so that the payment remains independent of θ_1 given o_1.

Deallocation In settings that have a null outcome and are *downward closed* (i.e., settings where a feasible outcome o remains feasible if o_i is replaced with the null outcome), we modify the function g to allocate the null outcome whenever the price function t_w creates an IR violation. This reduces ex post regret and in particular ensures ex post IR. On the other hand, the total value to the agents necessarily decreases under the modified allocation. In our experimental results, we refer to this as the *deallocation fix*.

3.5 Applying the Framework

In this section, we discuss the application of our framework to two domains: multi-minded combinatorial auctions and egalitarian welfare in the assignment problem.

3.5.1 Multi-Minded Combinatorial Auctions

A *combinatorial auction* allocates items $\{1,\ldots,r\}$ among n agents, such that each agent receives a possibly empty subset of the items. The outcome space Ω_i for agent i thus is the set of all subsets of the r items, and the type of agent i can be represented by a vector $\theta_i \in \Theta_i = \mathbb{R}^{2^r}$ that specifies its value for each possible bundle. The set of possible type profiles is then $\Theta = \mathbb{R}^{2^r n}$, and the value $v_i(\theta_i, o_i)$ of agent i for bundle o_i is equal to the entry in θ_i corresponding to o_i. We require that valuations are monotone, such that $v_i(\theta_i, o_i) \geq v_i(\theta_i, o_i')$ for all $o_i, o_i' \in \Omega_i$ with $o_i' \subseteq o_i$, and normalized such that $v_i(\theta_i, \emptyset) = 0$. Assuming agent symmetry and adopting the view of agent 1, the partial outcome rule $g_1 : \Theta \to \Omega_1$ specifies the bundle $g_1(\theta)$ allocated to agent 1; we require feasibility, so that no item is allocated more than once.

In a *multi-minded combinatorial auction* (multi-minded CA), each agent is interested in at most b bundles for some constant b. The special case where $b = 1$ is called a single-minded CA. In our framework, the restriction to multi-minded CAs leads to a number of computational advantages. First, valuation profiles and thus the training data can be represented in a compact way, by explicitly writing down the valuations for the constant number of bundles each agent is interested in. Second, inner products between valuation profiles, which are required to apply the kernel trick, can be computed in constant time.

Attribute Maps

To apply structural SVMs to multi-minded CAs, we need to specify an appropriate attribute map χ. In our experiments we use two attribute maps χ_1 and χ_2. The purpose of both attribute maps is to encode θ_{-1} and o_1. The first attribute map $\chi_1 : \Theta_{-1} \times \Omega_1 \to \mathbb{R}^{2^r(2^r(n-1))}$ achieves this by putting θ_{-1} into a particular position of a high-dimensional vector depending on o_1. The second attribute map $\chi_2 : \Theta_{-1} \times \Omega_1 \to \mathbb{R}^{2^r(n-1)}$ achieves this by restricting θ_{-1} to an outcome space without o_1. Formally,

$$\chi_1(\theta_{-1},o_1)=\begin{bmatrix}0\\ \cdots\\ 0\\ \theta_{-1}\\ 0\\ \cdots\\ 0\end{bmatrix}\begin{matrix}\left.\begin{matrix}\\ \\ \\ \end{matrix}\right\} dec(o_1)(2^r(n-1))\\ \\ \left.\begin{matrix}\\ \\ \\ \end{matrix}\right\}(2^r-dec(o_1)-1)(2^r(n-1))\end{matrix}, \quad \chi_2(\theta_{-1},o_1)=\begin{bmatrix}\theta_2\setminus o_1\\ \theta_3\setminus o_1\\ \cdots\\ \theta_n\setminus o_1\end{bmatrix}.$$

Here, $dec(o_1)=\sum_{j=1}^{r}2^{j-1}I_{j\in o_1}$ is a decimal index of bundle o_1, where $I_{j\in o_1}=1$ if $j\in o_1$ and $I_{j\in o_1}=0$ otherwise. Attribute map χ_1 thus stacks the vector θ_{-1}, which represents the valuations of all agents except agent 1, with zero vectors of the same dimension, where the position of θ_{-1} is determined by the index of bundle o_1. The resulting attribute vector is simple but potentially restrictive. It precludes two instances with different allocated bundles from sharing attributes, which provides an obstacle to generalization of the discriminant function across bundles. Attribute map χ_2 stacks vectors $\theta_i\setminus o_1$ for agents $i\neq 1$, where for each agent i the vector $\theta_i\setminus o_1$ is obtained from θ_i by setting the entries for all bundles that intersect with o_1 to 0. This captures the fact that agent i cannot be allocated any of the bundles that intersect with o_1 if o_1 is allocated to agent 1.[5]

Efficient Computation of Inner Products

For both χ_1 and χ_2, computing inner products reduces to the question of whether inner products between valuation profiles are efficiently computable. For χ_1, we have that

$$\left\langle \chi_1(\theta_{-1},o_1),\chi_1(\theta'_{-1},o'_1)\right\rangle = I_{o_1=o'_1}\sum_{i=2}^{n}\left\langle \theta_i,\theta'_i\right\rangle,$$

where indicator $I_{o_1=o'_1}=1$ if $o_1=o'_1$ and $I_{o_1=o'_1}=0$ otherwise. For χ_2,

$$\left\langle \chi_2(\theta_{-1},o_1),\chi_2(\theta'_{-1},o'_1)\right\rangle = \sum_{i=2}^{n}\left\langle \theta_i\setminus o_1,\theta'_i\setminus o_1\right\rangle.$$

We next develop efficient methods for computing the inner products $\langle\theta_i,\theta'_i\rangle$ on compactly represented valuation functions. The computation of $\langle\theta_i\setminus o_1,\theta'_i\setminus o_1\rangle$ can be done through similar methods.

In the single-minded setting, let θ_i correspond to a bundle $S_i\subseteq\{1,\ldots,r\}$ of items with value v_i, and θ'_i correspond to a set $S'_i\subseteq\{1,\ldots,r\}$ of items valued at v'_i.

[5]Both χ_1 and χ_2 are defined for a particular number of items and agents, and in our experiments we train a different classifier for each number of agents and items. In practice, one can pad out items and agents by setting bids to zero and train a single classifier.

Each set containing both S_i and S'_i contributes $v_i v'_i$ to $\theta_i^T \theta'_i$, while all other sets contribute 0. Since there are exactly $2^{r-|S_i \cup S'_i|}$ sets containing both S_i and S'_i, we have $\theta_i^T \theta'_i = v_i v'_i 2^{r-|S_i \cup S'_i|}$.

This is a special case of the formula for the multi-minded case.

Lemma 2. *Consider a multi-minded CA and two bid vectors x_1 and x'_1 corresponding to sets $S = \{S_1, \ldots, S_s\}$ and $S' = \{S'_1, \ldots, S'_t\}$, with associated values $v_1, \ldots, v_s$ and $v'_1, \ldots, v'_t$. Then,*

$$x_1^T x'_1 = \sum_{T \subseteq S, T' \subseteq S'} \Big((-1)^{|T|+|T'|} \cdot (\min_{S_i \in T} v_i) \cdot (\min_{S'_j \in T'} v'_j) \cdot 2^{r - |(\bigcup_{S_i \in T} S_i) \cup (\bigcup_{S'_j \in T'} S'_j)|} \Big). \quad (3.4)$$

Proof. The contribution of a particular bundle B of items to the inner product is $(\max_{S_i \in S, S_i \subseteq B} v_i) \cdot (\max_{S'_j \in S', S'_j \subseteq B} v'_j)$, and thus

$$x_1^T x'_1 = \sum_B \Big((\max_{\substack{S_i \in S \\ S_i \subseteq B}} v_i) \cdot (\max_{\substack{S'_j \in S' \\ S'_j \subseteq B}} v'_j) \Big).$$

By the maximum-minimums identity, which asserts that for any set $\{x_1, \ldots, x_n\}$ of n numbers, $\max\{x_1, \ldots, x_n\} = \sum_{Z \subseteq X} ((-1)^{|Z|+1} \cdot (\min_{x_i \in Z} x_i))$,

$$\max_{\substack{S_i \in S \\ S_i \subseteq B}} v_i = \sum_{\substack{T \subseteq S \\ \bigcup_{S_i \in T} S_i \subseteq B}} \Big((-1)^{|T|+1} \cdot (\min_{S_i \in T} v_i) \Big) \quad \text{and} \quad \max_{\substack{S'_j \in S' \\ S'_j \subseteq B}} v'_j = \sum_{\substack{T' \subseteq S' \\ \bigcup_{S'_j \in T'} S'_j \subseteq B}} \Big((-1)^{|T'|+1} \cdot (\min_{S'_j \in T'} v'_j) \Big).$$

The inner product can thus be written as

$$\theta_1^T \theta'_1 = \sum_B \sum_{\substack{T \subseteq S, T' \subseteq S' \\ \bigcup_{S_i \in T} S_i \subseteq B \\ \bigcup_{S'_j \in T'} S'_j \subseteq B}} \Big((-1)^{|T|+|T'|} \cdot (\min_{S_i \in T} v_i) \cdot (\min_{S'_j \in T'} v'_j) \Big).$$

Finally, for given $T \subseteq S$ and $T' \subseteq S'$, there exist exactly $2^{r - |(\bigcup_{S_i \in T} S_i) \cup (\bigcup_{S'_j \in T'} S'_j)|}$ bundles B such that $\bigcup_{S_i \in T} S_i \subseteq B$ and $\bigcup_{S'_j \in T'} S'_j \subseteq B$, and we obtain

$$\theta_1^T \theta'_1 = \sum_{T \subseteq S, T' \subseteq S'} \Big((-1)^{|T|+|T'|} \cdot (\min_{S_i \in T} v_i) \cdot (\min_{S'_j \in T'} v'_j) \cdot 2^{m - |(\bigcup_{S_i \in T} S_i) \cup (\bigcup_{S'_j \in T'} S'_j)|} \Big). \qquad \square$$

If S and S' have constant size, then the sum on the right hand side of (3.4) ranges over a constant number of sets and can be computed efficiently.

Dealing with an Exponentially Large Output Space

Recall that Training Problems 1 and 2 have constraints for every training example (θ^k, o_1^k) and every possible bundle of items $o_1 \in \Omega_1$, of which there are exponentially many in the number of items in the case of CAs. In lieu of an efficient separation oracle, a workaround exists when the discriminant function has additional structure, such that the induced payment weakly increases as items are added to a bundle. Given this *item monotonicity*, it would suffice to include constraints for bundles that have a strictly larger value to the agent than any of their respective subsets.

Still, it remains an open problem whether item monotonicity itself can be imposed on the hypothesis class with a small number of constraints.[6] An alternative is to optimistically assume item monotonicity, only including the constraints associated with bundles that are explicit in agent valuations. The baseline experimental results in Section 3.6 do not assume item monotonicity and instead use a separation oracle that iterates over all possible bundles $o_1 \in \Omega_1$. We also present results which test the idea of optimistically assuming item monotonicity, and while there is a degradation in performance, results are mostly comparable.

3.5.2 The Assignment Problem

In the *assignment problem*, we are given a set of n agents and a set $\{1,\dots,n\}$ of items, and wish to assign each item to exactly one agent. The outcome space of agent i is thus $\Omega_i = \{1,\dots,n\}$, and its type can be represented by a vector $\theta_i \in \Theta_i = \mathbb{R}^n$. The set of possible type profiles is then $\Theta = \mathbb{R}^{n^2}$. We consider an outcome rule that maximizes *egalitarian welfare* in a lexicographic manner: first, the minimum value of any agent is maximized; if more than one outcome achieves the minimum, the second lowest value is maximized, and so forth. This outcome rule can be computed by solving a sequence of integer programs. As before, we assume agent symmetry and adopt the view of agent 1.

Attribute Map

We need to define an attribute map $\chi_3 : \mathbb{R}^{n^2-n} \times \mathbb{N} \to \mathbb{R}^s$, where the first argument is the type profile of all agents but agent 1, the second argument is the item assigned to agent 1, and s is a dimension of our choosing. A natural choice for χ_3 is to set

$$\chi_3(\theta_{-1}, j) = (\theta_2[-j], \theta_3[-j], \dots, \theta_n[-j]) \in \mathbb{R}^{(n-1)^2},$$

where $\theta_i[-j]$ denotes the vector obtained from θ_i by removing the jth entry. The attribute map thus reflects the agents' values for all items except item j, capturing the fact that the item assigned to agent 1 cannot be assigned to any other agent.

Note that in the case of the assignment problem neither the computation of inner products nor searching the outcome space poses a computational problem.

[6] For polynomial kernels and certain attribute maps, a possible sufficient condition for item monotonicity is to force the weights w_{-1} to be negative. However, as with the discussion of enforcing $w_1 > 0$ directly, these weight constraints do not dualize conveniently and results in the dual formulation no longer operate on inner products $\langle \psi'(\theta_{-1}, o_1), \psi'(\theta'_{-1}, o'_1) \rangle$. As a result, we would be forced to work in the primal, and incur extra computational overhead that increases polynomially with the kernel degree d. We have performed some preliminary experiments with polynomial kernels, but we have not looked into reformulating the primal to enforce item monotonicity.

3.6 Experimental Evaluation

We perform a series of experiments to test our theoretical framework. To run our experiments, we use the *SVM^{struct}package* [12], which allows for the use of custom kernel functions, attribute maps, and separation oracles.

3.6.1 Setup

We begin by briefly discussing our experimental methodology, performance metrics, and optimizations used to speed up the experiments.

Methodology

For each of the settings we consider, we generate three data sets: a *training set*, a *validation set*, and a *test set*. The training set is used as input to Training Problem 2, which in turn yields classifiers h_w and corresponding payment rules p_w. For each choice of the parameter C of Training Problem 2, and the parameter γ if the RBF kernel is used, a classifier h_w is learned based on the training set and evaluated based on the validation set. The classifier with the highest accuracy on the validation set is then chosen and evaluated on the test set. During training, we take the perspective of agent 1, so a training set size of ℓ means that we train an SVM on ℓ examples. Once a partial outcome rule has been learned, however, it can be used to infer payments for all agents. We exploit this fact during testing, and report performance metrics across all agents for a given instance in the test set.

Metrics

We employ three metrics to measure the performance of the learned classifiers. These metrics are computed over the test set $\{(\theta^k, o^k)\}_{k=1}^{\ell}$.

Classification accuracy *Classification accuracy* measures the accuracy of the trained classifier in predicting the outcome. Each instance of the ℓ instances has n agents, so in total we measure accuracy over $n\ell$ instances:[7]

$$\textit{accuracy} = 100 \cdot \frac{\sum_{k=1}^{\ell} \sum_{i=1}^{n} I(h_w(\theta_i, \theta_{-i}) = o_i^k))}{n\ell}.$$

[7] For a given instance θ, there are actually many ways to choose (θ_i, θ_{-i}) depending on the ordering of all agents but agent i. We discuss a technique we refer to as sorting in Section 3.6.1, which will choose a particular ordering. When this technique is not used, for example in our experiments for the assignment problem, we simply fix an ordering of the other agents for each agent i and use the same ordering across all instances.

Ex post regret We measure *ex post regret* by summing over the ex post regret experienced by all agents in each of the ℓ instances in the dataset, i.e.,

$$regret = \frac{\sum_{k=1}^{\ell}\sum_{i=1}^{n} rgt_i(\theta_i^k, \theta_{-i}^k)}{n\ell}.$$

Individual rationality violation This metric measures the fraction of *individual rationality violation* across all agents:

$$ir\text{-}violation = \frac{\sum_{k=1}^{\ell}\sum_{i=1}^{n} I(irv_i(\theta_i, \theta_{-i}) > 0)}{n\ell}.$$

Optimizations

In the case of multi-minded CAs we map the inputs θ_{-1} into a smaller space, which allows us to learn more effectively with smaller amounts of data.[8] We use *instance-based normalization*, which normalizes the values in θ_{-1} by the highest observed value and then rescales the computed payment appropriately, and *sorting*, which orders agents based on bid values.

Instance-Based Normalization The first technique we use is *instance-based normalization*. Before passing examples θ to the learning algorithm or learned classifier, they are normalized by a positive multiplier so that the value of the highest bid by agents other than agent 1 is exactly 1, before passing it to the learning algorithm or classifier. The values and the solution are then transformed back to the original scale before computing the payment rule p_w. This technique leverages the observation that agent 1's allocation depends on the relative values of the other agent's reports (scaling all reports by a factor should not affect the outcome chosen).

Sorting The second technique we use is *sorting*. With sorting, instead of choosing an arbitrary ordering of agents in θ_{-i}, we choose a specific ordering based on the maximum value the agent reports. In the single-item setting, this amounts to ordering agents by their value. In the multi-minded CA setting, agents are ordered by the value they report for their most desired bundle. The intuition behind sorting is that we can again decrease the space of possible θ_{-i} reports the learner sees and learn more quickly. In the single-item case, we know that the second price payment rule only depends on the maximum value across all other agents, and sorting places this value in the first coordinate of θ_{-i}.

[8] The barrier to using more data is not the availability of the data itself, but the time required for training, because training time scales quadratically in the size of the training set due to the use of non-linear kernels.

3.6.2 Single-Item Auction

As a sanity check, we perform experiments on the single-item auction with the optimal outcome rule, where the agent with the highest bid receives the item. For this outcome rule g we know that the associated payment rule p that makes (g, p) strategyproof is the second price payment rule.

For our experiments we draw the values of the agents from the uniform distribution D on $[0, 1]$. We use a training set size of 300 and validation and test set sizes of 1000. We use the attribute maps χ_1 and χ_2, which can be applied to this setting because single-item auctions are a special case of multi-minded CAs. In particular, letting z be the 0 vector of dimension $n-1$, $\chi_1(\theta_{-1}, o_1) = (\theta_{-1}, z)$ if $o_1 = \emptyset$ and $\chi_1(\theta_{-1}, o_1) = (z, \theta_{-1})$ if $o_1 = \{1\}$ and $\chi_2(\theta_{-1}, o_1) = \theta_{-1}$ if $o_1 = \emptyset$ and $\chi_2(\theta_{-1}, o_1) = z$ if $o_1 = \{1\}$. We use the regularization parameter $C \in \{10^4, 10^5\}$ and the RBF kernel with parameter $\gamma \in \{0.01, 0.1, 1\}$.

Table 3.1 and Figure 3.1 show that for both attribute maps we obtain excellent accuracy and very close approximation to the second price payment rule. This shows that the framework is able to automatically learn the payment rule of Vickrey's auction.

n	accuracy		regret		ir-violation	
	χ_1	χ_2	χ_1	χ_2	χ_1	χ_2
2	99.7	93.1	0.000	0.003	0.00	0.07
3	98.7	97.6	0.000	0.000	0.01	0.00
4	98.4	99.1	0.000	0.000	0.00	0.01
5	97.3	96.6	0.001	0.001	0.02	0.00
6	97.6	97.4	0.000	0.001	0.00	0.02

Table 3.1: Basic performance metrics for single-item auction.

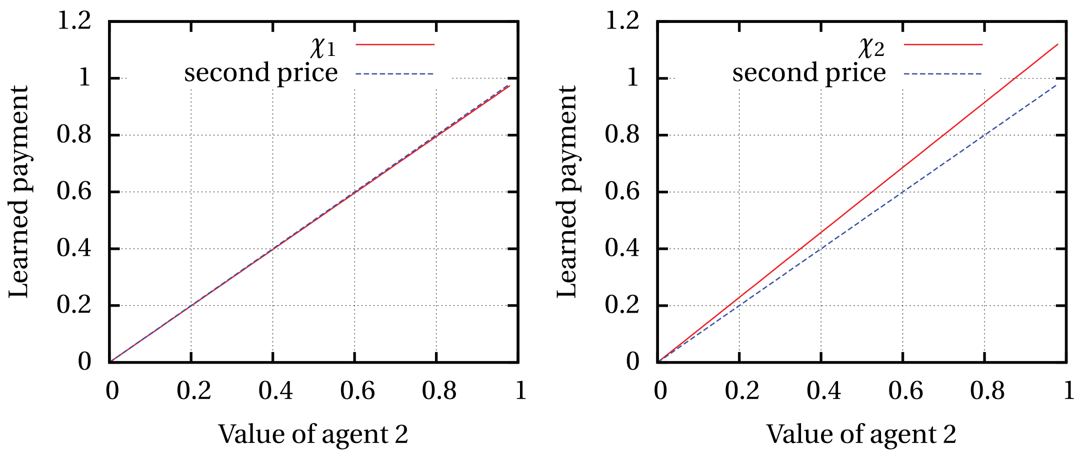

Figure 3.1: Learned payment rule vs. second price payment rule in a single-item auction with two agents for χ_1 (left) and χ_2 (right).

3.6.3 Multi-Minded Combinatorial Auctions

Type Distribution

Recall that in a multi-minded setting, there are r items, and each agent is interested in exactly b bundles. For each bundle, we use the following procedure (inspired by Sandholm's *decay distribution* for the single-minded setting [25]) to determine which items are included in the bundle. We first assign an item to the bundle uniformly at random. Then with probability α, we add another random item (chosen uniformly from the remaining items), and with probability $(1-\alpha)$ we stop. We continue this procedure until we stop or have exhausted the items. We use $\alpha = 0.75$ to be consistent with [25], as they report that the winner determination problem (finding the feasible allocation that maximizes total value) is difficult for this setting of α.

Once the bundle identities have been determined, we sample values for these bundles. Let c be an r-dimensional vector with entries chosen uniformly from $(0,1]$. For each agent i, let d_i be an r-dimensional vector with entries chosen uniformly from $(0,1]$. Each entry of c denotes the common value of a specific item, while each entry of d_i denotes the private value of a specific item for agent i. The value of bundle S_{ij} is then given by

$$v_{ij} = \min_{S_{ij'} \leq S_{ij}} \left(\frac{\langle S_{ij'}, \beta c + (1-\beta) d_i \rangle}{r} \right)^{\zeta}$$

for parameters $\beta \in [0,1]$ and ζ. The inner product in the numerator corresponds to a sum over values of items, where common and private values for each item are respectively weighted with β and $(1-\beta)$. The denominator normalizes all valuations to the interval $(0,1]$. Parameter ζ controls the degree of *complementarity* among items: $\zeta > 1$ implies that goods are complements, whereas $\zeta < 1$ means that goods are substitutes. Choosing the minimum over bundles $S_{ij'}$ contained in S_{ij} finally ensures that the resulting valuations are monotonic.

Outcome Rules

We use two outcome rules in our experiments. The *optimal outcome rule* g_{opt} assigns the bundles such that the social welfare is maximized. In the case of the optimal outcome rule there is a payment rule p_{vcg} that makes the mechanism (g_{opt}, p_{vcg}) strategyproof. From agent 1's perspective this payment rule is given by

$$p_{vcg,1}(\theta) = \left(\max_{o \in \Omega} \sum_{i \neq 1} v_i(\theta_i, o_i) \right) - \sum_{i \neq 1} v_i(\theta_i, g_{opt,i}(\theta)).$$

The second outcome rule with which we experiment is a generalization of the *greedy outcome rule* for single-minded CA Lehmann et al. [16]. Our generalization of the greedy outcome rule is as follows. Let θ be the agent valuations and $o_i(j)$ denote the j-th bundle desired by agent i. For each bundle $o_i(j)$, assign a score $v_i(\theta_i, o_i(j))/\sqrt{|o_i(j)|}$, where $|o_i(j)|$ indicates the total items in bundle $o_i(j)$. The greedy outcome rule orders the desired bundles by this score, and takes the bundle $o_i(j)$ with the next highest score as long as agent i has not already been allocated a bundle and $o_i(j)$ does not contain any items already allocated. While this greedy outcome rule has an associated payment rule that makes it strategyproof in the single-minded case, it is not implementable in the multi-minded case as the following example shows.

Example 1. Consider a setting with a single agent and four items.

If the valuations θ_1 of the agent are

$$v_1(\theta_1, o_1) = \begin{cases} 20 & \text{if } o_1 = \{1,2,3,4\} \\ 12 & \text{if } 1 \in o_1 \text{ and } j \notin o_1 \text{ for some } j \in \{2,3,4\}, \text{ and} \\ 0 & \text{else} \end{cases}$$

then the allocation is $\{1\}$.

If the valuations are θ_1' such that

$$v_1(\theta_1', o_1) = \begin{cases} 12 & \text{if } o_1 = \{1,2,3,4\} \\ 5 & \text{if } 1 \in o_1 \text{ and } j \notin o_1 \text{ for some } j \in \{2,3,4\}, \text{ and} \\ 0 & \text{else} \end{cases}$$

then the allocation is $\{1,2,3,4\}$.

We have $v_1(\theta_1', \{1,2,3,4\}) - v_1(\theta_1', \{1\}) < v_1(\theta_1, \{1,2,3,4\}) - v_1(\theta_1, \{1\})$ contradicting weak monotonicity.

Description of Experiments

We experiment with training sets of sizes 100, 300, and 500, and validation and test sets of size 1000. All experiments we report on are for a setting with 5 agents, 5 items, and 3 bundles per agent, and use $\beta = 0.5$, the RBF kernel, and parameters $C \in \{10^4, 10^5\}$ and $\gamma \in \{0.01, 0.1, 1\}$.

Basic Results

Table 3.2 presents the basic results for multi-minded CAs with a training set size of 500 for both the optimal and the greedy outcome rule. For both outcome rules, we present the results for p_{vcg} as a baseline. Because p_{vcg} is the strategyproof payment rule for the optimal outcome rule, p_{vcg} always has accuracy 100, regret 0, and IR violation 0 for the optimal outcome rule.

Across all instances, as expected, accuracy is negatively correlated with regret and ex post IR violation. The degree of complementarity between items, ζ, as well as the outcome rule chosen, has a major effect on the results. Instances with low complementarity ($\zeta = 0.5$) yield

		Optimal outcome rule									Greedy outcome rule								
		accuracy			regret			ir-violation			accuracy			regret			ir-violation		
n	ζ	p_{vcg}	χ_1	χ_2	p_{vcg}	χ_1	χ_2	p_{vcg}	χ_1	χ_2	p_{vcg}	χ_1	χ_2	p_{vcg}	χ_1	χ_2	p_{vcg}	χ_1	χ_2
2	0.5	100	70.7	91.9	0	0.014	0.002	0	0.06	0.03	50.9	59.1	40.6	0.079	0.030	0.172	0.22	0.12	0.33
3	0.5	100	54.5	75.4	0	0.037	0.017	0	0.19	0.10	55.4	57.9	54.7	0.070	0.030	0.088	0.18	0.21	0.36
4	0.5	100	53.8	67.7	0	0.042	0.031	0	0.22	0.18	61.1	58.2	57.9	0.056	0.033	0.056	0.14	0.20	0.31
5	0.5	100	15.8	67.0	0	0.133	0.032	0	0.26	0.19	64.9	61.3	63.0	0.048	0.027	0.042	0.13	0.19	0.24
6	0.5	100	61.1	68.2	0	0.037	0.032	0	0.22	0.20	66.6	63.8	63.8	0.041	0.034	0.045	0.12	0.20	0.24
2	1.0	100	84.5	93.4	0	0.008	0.001	0	0.08	0.02	87.8	86.6	84.0	0.007	0.005	0.008	0.04	0.06	0.09
3	1.0	100	77.1	83.5	0	0.012	0.005	0	0.13	0.09	85.3	86.7	85.7	0.006	0.006	0.006	0.04	0.07	0.05
4	1.0	100	74.6	81.1	0	0.014	0.009	0	0.16	0.12	82.4	86.5	84.2	0.006	0.006	0.007	0.05	0.08	0.08
5	1.0	100	73.4	77.4	0	0.018	0.011	0	0.19	0.12	82.7	85.8	84.9	0.007	0.009	0.009	0.04	0.10	0.10
6	1.0	100	75.0	77.7	0	0.020	0.013	0	0.20	0.16	80.0	87.4	88.1	0.006	0.007	0.005	0.04	0.08	0.07
2	1.5	100	91.5	96.9	0	0.004	0.000	0	0.06	0.02	94.7	91.1	91.7	0.002	0.002	0.002	0.02	0.04	0.04
3	1.5	100	91.0	93.4	0	0.004	0.001	0	0.05	0.03	97.1	92.8	93.2	0.001	0.002	0.001	0.01	0.02	0.04
4	1.5	100	92.5	94.2	0	0.003	0.001	0	0.03	0.04	96.4	91.5	92.1	0.001	0.003	0.002	0.02	0.07	0.07
5	1.5	100	91.7	93.9	0	0.004	0.002	0	0.06	0.03	97.5	90.5	91.4	0.001	0.004	0.002	0.01	0.06	0.04
6	1.5	100	91.9	93.7	0	0.003	0.001	0	0.05	0.04	98.4	92.2	92.8	0.000	0.003	0.002	0.01	0.06	0.06

Table 3.2: Basic performance metrics for multi-minded CA.

payment rules with higher regret, and χ_1 performs better on the greedy outcome rule while χ_2 performs better on the optimal outcome rule. For high complementarity between items the greedy outcome tends to allocate all items to a single agent, and the learned price function sets high prices for small bundles to capture this property. For low complementarity the allocation tends to be split and less predictable. Still, the best classifiers achieve average ex post regret of less than 0.032 (for values normalized to [0,1]) even though the corresponding prediction accuracy can be as low as 67%. For the greedy outcome rule, the performance of p_{vcg} is comparable for $\zeta \in \{1.0, 1.5\}$ but worse than the payment rule learned in our framework in the case of $\zeta = 0.5$, where the greedy outcome rule becomes less optimal.

Effect of Training Set Size

Table 3.3 charts performance as the training set size is varied for the greedy outcome rule. While training data is readily available (we can simply sample from D and run the outcome rule g), training time becomes prohibitive for larger training set sizes. Table 3.3 shows that regret decreases with larger training sets, and for a training set size of 500, the best of χ_1 and χ_2 outperforms p_{vcg} for $\zeta = 0.5$ and is comparable to p_{vcg} for $\zeta \in \{1.0, 1.5\}$.

Individual Rationality Violations

Table 3.4 summarizes our results regarding the various fixes to IR violations for the greedy outcome rule with attribute map χ_2 and the particularly challenging case $\zeta = 0.5$. Each row corresponds to a different payment offset and each column represents a different null loss. The extent of IR violation decreases with larger payment offset and null loss. Regret tends to move

n	ζ	accuracy	100		300		500		regret	100		300		500	
		p_{vcg}	χ_1	χ_2	χ_1	χ_2	χ_1	χ_2	p_{vcg}	χ_1	χ_2	χ_1	χ_2	χ_1	χ_2
2	0.5	50.9	54.3	48.2	57.0	46.9	59.1	40.6	0.079	0.045	0.195	0.032	0.098	0.030	0.172
3	0.5	55.4	50.1	49.8	55.7	54.4	57.9	54.7	0.070	0.054	0.078	0.038	0.082	0.030	0.088
4	0.5	61.1	53.4	56.2	56.4	58.5	58.2	57.9	0.056	0.050	0.059	0.040	0.061	0.033	0.056
5	0.5	64.9	14.2	57.9	61.0	61.8	61.3	63.0	0.048	0.173	0.064	0.038	0.048	0.027	0.042
6	0.5	66.6	58.4	58.8	62.2	63.9	63.8	63.8	0.041	0.039	0.059	0.037	0.049	0.034	0.045
2	1.0	87.8	80.7	80.5	84.4	84.1	86.6	84.0	0.007	0.010	0.010	0.009	0.008	0.005	0.008
3	1.0	85.3	74.9	78.0	83.0	80.6	86.7	85.7	0.006	0.020	0.011	0.009	0.009	0.006	0.006
4	1.0	82.4	78.5	80.1	84.2	83.1	86.5	84.2	0.006	0.015	0.014	0.008	0.009	0.006	0.007
5	1.0	82.7	81.0	81.8	84.3	84.3	85.8	84.9	0.007	0.020	0.014	0.010	0.009	0.009	0.009
6	1.0	80.0	81.8	83.7	87.6	88.3	87.4	88.1	0.006	0.062	0.018	0.008	0.005	0.007	0.005
2	1.5	94.7	83.3	88.1	89.3	89.8	91.1	91.7	0.002	0.008	0.003	0.003	0.002	0.002	0.002
3	1.5	97.1	86.9	87.6	90.3	91.5	92.8	93.2	0.001	0.005	0.004	0.003	0.002	0.002	0.001
4	1.5	96.4	88.4	90.7	89.3	90.8	91.5	92.1	0.001	0.005	0.003	0.004	0.003	0.003	0.002
5	1.5	97.5	87.2	88.5	91.4	90.5	90.5	91.4	0.001	0.006	0.004	0.003	0.003	0.004	0.002
6	1.5	98.4	86.3	86.8	91.4	92.5	92.2	92.8	0.000	0.011	0.007	0.004	0.002	0.003	0.002

Table 3.3: Effect of training set size on accuracy and regret in a multi-minded CA with a greedy outcome rule.

payment	accuracy			regret			ir-violation			ir-fix-welfare-avg		
offset	0.5	1.0	1.5	0.5	1.0	1.5	0.5	1.0	1.5	0.5	1.0	1.5
0	59.7	61.8	61.7	0.065	0.048	0.042	0.35	0.26	0.21	0.27	0.43	0.52
0.05	61.7	61.2	60.1	0.054	0.045	0.044	0.29	0.20	0.15	0.37	0.54	0.65
0.10	62.1	59.3	56.7	0.048	0.047	0.051	0.23	0.14	0.10	0.48	0.66	0.75
0.15	60.4	55.1	52.2	0.047	0.055	0.064	0.17	0.10	0.06	0.59	0.75	0.84
0.20	57.8	51.7	48.5	0.052	0.067	0.079	0.12	0.06	0.03	0.70	0.83	0.90
0.25	54.3	47.7	44.3	0.061	0.082	0.096	0.08	0.03	0.02	0.79	0.89	0.93

Table 3.4: Impact of payment offset and null loss fix in a multi-minded CA with a greedy outcome rule.

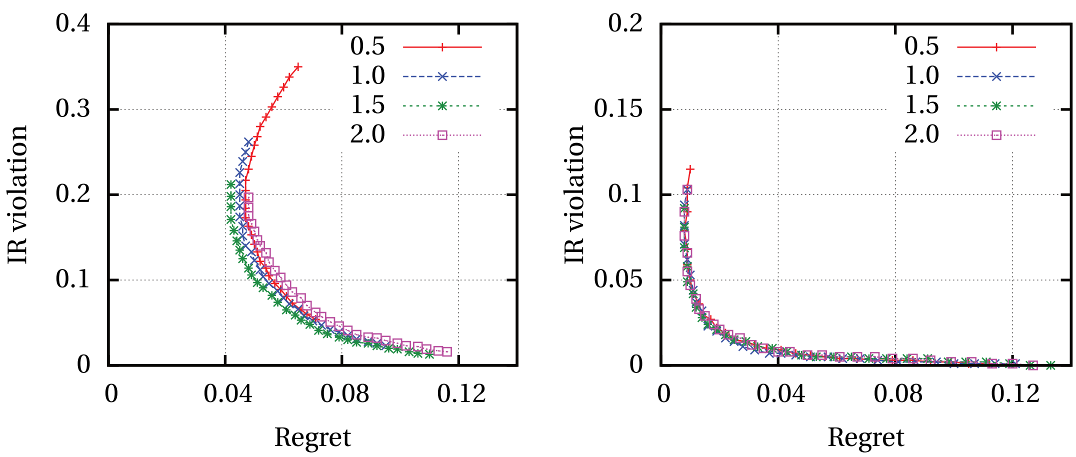

Figure 3.2: Impact of payment offset and null loss fix in a multi-minded CA with a greedy outcome rule for $\zeta = 0.5$ (left) and $\zeta = 1.0$ (right).

n	ζ	accuracy		regret		ir-violation	
		χ_2	χ_2 (i-mon)	χ_2	χ_2 (i-mon)	χ_2	χ_2 (i-mon)
2	0.5	46.9	46.3	0.098	0.232	0.28	0.38
3	0.5	54.4	8.6	0.082	0.465	0.33	0.06
4	0.5	58.5	48.2	0.061	0.811	0.31	0.25
5	0.5	61.8	57.0	0.048	0.136	0.26	0.26
6	0.5	63.9	61.3	0.049	0.078	0.25	0.20
2	1.0	84.1	82.2	0.008	0.010	0.06	0.08
3	1.0	80.6	80.1	0.009	0.010	0.10	0.09
4	1.0	83.1	79.7	0.009	0.012	0.11	0.11
5	1.0	84.3	77.2	0.009	0.020	0.10	0.11
6	1.0	88.3	83.9	0.005	0.013	0.08	0.11
2	1.5	89.8	89.1	0.002	0.003	0.03	0.06
3	1.5	91.5	91.3	0.002	0.003	0.04	0.04
4	1.5	90.8	89.7	0.003	0.003	0.06	0.06
5	1.5	90.5	87.3	0.003	0.005	0.04	0.05
6	1.5	92.5	70.8	0.002	0.081	0.06	0.17

Table 3.5: Comparison of performance with and without optimistically assuming item monotonicity in a multi-minded CA with a greedy outcome rule.

in the opposite direction, but there are cases where IR violation and regret both decrease. The three rightmost columns of Table 3.4 list the average ratio between welfare after and before the deallocation fix, across the instances in the test set. With a payment offset of 0, a large welfare hit is incurred if we deallocate agents with IR violations. However, this penalty decreases with increasing payment offsets and increasing null loss. At the most extreme payment offset and null loss adjustment, the IR violation is as low as 2% and the deallocation fix incurs a welfare loss of only 7%.

Figure 3.2 shows a graphical representation of the impact of payment offsets and null losses for the greedy outcome rule with attribute map χ_2. The left plot is for $\zeta = 0.5$ and the right plot is for $\zeta = 1$. Each line in the plot corresponds to a payment rule learned with a different null loss, and each point on a line corresponds to a different payment offset. The payment offset is zero for the top-most point on each line, and equal to 0.29 for the lowest point on each line. Increasing the payment offset always decreases the rate of IR violation, but may decrease or increase regret. Increasing null loss lowers the top-most point on a given line, but arbitrarily increasing null loss can be harmful. Indeed, in the figure on the left, a null loss of 1.5 results in a slightly higher top-most point but significantly lower regret at this top-most point compared to a null loss of 2.0. It is also interesting to note that these adjustments have much more impact for $\zeta = 0.5$.

Item Monotonicity

Table 3.5 presents a comparison of a payment rule learned with explicit enumeration of all bundle constraints (the default that we have been using for our other results) and a payment rule learned by optimistically assuming item monotonicity (see Section 3.5.1). Performance is

n	accuracy				regret				ir-violation			
	vcg	tot-vcg	eg-vcg	p_w	vcg	tot-vcg	eg-vcg	p_w	vcg	tot-vcg	eg-vcg	p_w
2	64.3	67.5	67.5	89.0	0.018	0.015	0.015	0.023	0.03	0.01	0.01	0.03
3	48.0	52.1	42.5	77.9	0.070	0.077	0.127	0.041	0.06	0.07	0.03	0.04
4	40.6	43.1	30.8	71.0	0.111	0.123	0.199	0.054	0.07	0.09	0.03	0.02
5	32.4	35.3	24.5	63.9	0.157	0.169	0.254	0.071	0.10	0.12	0.03	0.01
6	27.1	29.9	20.0	59.0	0.189	0.208	0.290	0.074	0.10	0.13	0.03	0.01

Table 3.6: Basic Performance metrics for assignment problem with egalitarian outcome rule

affected when we drop constraints and optimistically assume item monotonicity, although the effects are small for $\zeta \in \{1.0, 1.5\}$ and larger for $\zeta = 0.5$. Because item monotonicity allows for the training problem to be succinctly specified, we may be able to train on more data, and this seems a very promising avenue for further consideration (perhaps coupled with heuristic methods to add additional constraints to the training problem).

3.6.4 The Assignment Problem

In the assignment problem, agents' values for the items are sampled uniformly and independently from $[0,1]$. We use a training set of size 600, validation and test sets of size 1000, regularization parameter $C \in \{10, 1000, 100000\}$, and the RBF kernel with parameter $\gamma \in \{0.1, 0.5, 1.0\}$.

The performance of the learned payment rules is compared to that of three VCG-based payment rules. Let W be the total welfare of all agents other than i under the outcome chosen by g, and W_{eg} be the minimum value any agent other than i receives under this outcome. We then consider the following payment rules: (1) the *vcg* payment rule, where agent i pays the difference between the maximum total welfare of the other agents under any allocation and W; (2) the *tot-vcg* payment rule, where agent i pays the difference between the total welfare of the other agents under the allocation maximizing egalitarian welfare and W; and (3) the *eg-vcg* payment rule, where agent i pays the difference between the minimum value of any agent under the allocation maximizing egalitarian welfare and W_{eg}.

The results for attribute map χ_3 are shown in Table 3.6. We see that the learned payment rule p_w yields significantly lower regret than any of the VCG-based payment rules, and average ex post regret less than 0.074 for values normalized to $[0,1]$. Since we are not maximizing the sum of values of the agents, it is not very surprising that VCG-based payment rules perform rather poorly. The learned payment rule p_w can adjust to the outcome rule, and also achieves a low fraction of ex post IR violation of at most 3%.

3.7 Conclusion and Future Work

We have introduced a new paradigm for computational mechanism design in which statistical machine learning is adopted to design payment rules for given algorithmically specified outcome rules, and have shown encouraging experimental results. Future directions of interest include (1) an alternative formulation of the problem as a regression rather than classification problem, (2) constraints on properties of the learned payment rule, concerning for example the core or budgets, (3) methods that learn classifiers more likely to induce feasible outcome rules, so that these learned outcome rules can be used, (4) optimistically assuming item monotonicity and dropping constraints implied by it, thereby allowing for better scaling of training time with training set size at the expense of optimizing against a subset of the full constraints in the training problem, and (5) an investigation of the extent to which alternative goals such as regret percentiles or *interim* regret can be achieved through machine learning.

Bibliography

1. I. Ashlagi, M. Braverman, A. Hassidim, and D. Monderer. Monotonicity and implementability. *Econometrica*, 78(5):1749–1772, 2010.

2. L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and P. Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, 2006.

3. X. Bei and Z. Huang. Bayesian incentive compatibility via fractional assignments. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 720–733, 2011.

4. E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119:1061–1103, 2011.

5. Y. Cai, C. Daskalakis, and S. M. Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, page Forthcoming, 2012.

6. G. Carroll. A quantitative approach to incentives: Application to voting rules. Technical report, MIT, 2011.

7. V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence*, pages 103–110, 2002.

8. R. Day and P. Milgrom. Core-selecting package auctions. *International Journal of Game Theory*, 36(3–4):393–407, 2008.

9. A. Erdil and P. Klemperer. A new payment rule for core-selecting package auctions. *Journal of the European Economic Association*, 8(2–3):537–547, 2010.

10. M. Guo and V. Conitzer. Computationally feasible automated mechanism design: General approach and case studies. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 2010.

11. J. D. Hartline, R. Kleinberg, and A. Malekian. Bayesian incentive compatibility via matchings. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 734–747, 2011.

12. T. Joachims, T. Finley, and C.-N. J. Yu. Cutting-plane training of structural SVMs. *Machine Learning*, 77(1):27–59, 2009.

13. S. Lahaie. A kernel method for market clearing. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 208–213, 2009.

14. S. Lahaie. Stability and incentive compatibility in a kernel-based combinatorial auction. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 811–816, 2010.

15. R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. In *Proceedings of the 46th Symposium on Foundations of Computer Science*, pages 595–604, 2005.

16. D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49:577–602, 2002.

17. B. Lubin. *Combinatorial Markets in Theory and Practice: Mitigating Incentives and Facilitating Elicitation*. PhD thesis, Department of Computer Science, Harvard University, 2010.

18. B. Lubin and D. C. Parkes. Quantifying the strategyproofness of mechanisms via metrics on payoff distributions. In *Proceedings of the 25th Annual Conference on Uncertainty in Artificial Intelligence*, pages 349–358, 2009.

19. B. Lubin and D. C. Parkes. Approximate strategyproofness. *Current Science*, 103(9):1021–1032, 2012.

20. R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

21. D. C. Parkes, J. Kalagnanam, and M. Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 1161–1168, 2001.

22. P. Pathak and T. Sönmez. Comparing mechanisms by their vulnerability to manipulation. *American Economic Review*, 103:80–106, 2013.

23. B. Rastegari, A. Condon, and K. Leyton-Brown. Revenue monotonicity in deterministic, dominant-strategy combinatorial auctions. *Artificial Intelligence*, 175:441–456, 2011.

24. M. Saks and L. Yu. Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 286–293, 2005.

25. T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

26. I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6: 1453–1484, 2005.

Concluding Remarks

In this thesis we addressed three challenges in algorithmic mechanism design, which combines the emphasis on computational complexity of algorithm design with the emphasis on strategic behavior of mechanism design. The challenges that we addressed were the design of expressive mechanisms, the analysis of simplicity-expressiveness tradeoffs, and the design of approximately strategyproof mechanisms.

For the first challenge we considered the domain of multi-item auctions with unit demand and presented the most expressive polynomial-time mechanism for this setting known to date that is incentive compatible for non-degenerate inputs. An interesting direction for future work would be to *push the expressiveness frontier even further*, i.e., find even more expressive mechanisms that can be executed in polynomial time and have good incentive properties. One possible avenue would be to consider more general non-linear and discontinuous utility functions. Another possibility would be to consider more general domains, for example, one-to-many or many-to-many matchings. In both cases a promising approach seems to be to relax the requirements, either by replacing exact bidder optimality or exact envy freeness with approximate bidder optimality or approximate envy freeness or by replacing exact incentive compatibility with approximate incentive compatibility.

For the second challenge we considered simplified mechanisms that result from restricting the message space of a reference mechanism, and analyzed the impact that these restrictions have on the set of equilibria as a whole. An interesting target for future work would be to *replace the revelation principle*, with its implied focus on direct-revelation mechanisms and blindness towards computational aspects, *with a more practical simplification principle* that allows to assess the advantages and disadvantages of simplification taking computational aspects into account. A particularly interesting question is how the computational burden should be shared between the mechanism and the agents. As a concrete example consider a domain in which computing the outcome for the true preferences is a computationally hard problem. The designer could sidestep the computational hardness on the mechanism's side by a harsh enough restriction of the message space, but what would the computational implications of such a harsh restriction be for the agents? Specifically, if the decision is between computational tractability of the mechanism or computational tractability of the strategic reasoning required by the agents, which is to be preferred?

For the third challenge we adopted expected ex post regret as a quantifiable target of approximate strategyproofness, and presented a framework that given an algorithmically specified outcome rule automatically finds a payment rule that makes the resulting mechanism maximally strategyproof. An interesting direction for future work, apart from the obvious improvements, such as trying to make the method more scalable, would be to *exploit the structural similarity between discriminant-based classification and strategyproof mechanism design in a different way*. One could, for example, imagine to use the learned discriminant-based classifier to derive both an outcome rule and a payment rule. The resulting mechanism would be guaranteed to be strategyproof, but the learned outcome rule would only be close to the original outcome rule. Similarly, one could imagine to adopt the framework to mechanism design without money. The idea would again be to use the learned outcome rule instead of the original outcome rule. In both cases the main obstacle seems to be that the learned outcome rule, at least if it is learned on an per-agent basis, may not be feasible.

Bibliographic Note

This thesis is based on the following publications. Each of the publications is the product of joint work with the co-authors listed.

1. Paul Dütting, Monika Henzinger, and Ingmar Weber. An expressive mechanism for auctions on the web. In *Proceedings of the 20th International World Wide Web Conference*, pages 127-136, 2011.

2. Paul Dütting, Felix Fischer, and David C. Parkes. Simplicity-expressiveness tradeoffs in mechanism design. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, pages 341-350, 2011.

3. Paul Dütting, Felix Fischer, Pichayut Jirapinyo, John K. Lai, Benjamin Lubin, and David C. Parkes. Payment rules through discriminant-based classifiers. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 477-494, 2012.

Curriculum Vitae

Paul Dütting

March 28, 2013

Coordinates

EPFL IC LTAA		Phone:	+41-21-6931311
BC 100 (Bâtiment BC)		Fax:	+41-21-6931310
Station 14		Email:	paul.duetting@epfl.ch
CH - 1015 Lausanne		Web:	http://www.paulduetting.com/

Personal

Date of birth:	May 28, 1983
Place of birth:	Yaoundé, Cameroon
Citizenship:	German

Research Interests

Computer Science, Economics; Algorithmic Game Theory, Algorithmic Mechanism Design

Education

September 2008–today	**EPFL (Lausanne, Switzerland)** Ph.D. in Computer Science, March 2013 Advisors: Monika Henzinger, Boi Faltings
October 2003–May 2008	**KIT (Karlsruhe, Germany)** M.Sc. in Computer Science, "with distinction", May 2008 Advisor: Dorothea Wagner

Academic Employment

September 2008–today	**EPFL (Lausanne, Switzerland)** Position: Ph.D. student Advisors: Monika Henzinger, Boi Faltings

Awards and Fellowships

May 2013–April 2014	SNF Postdoctoral Fellowship
June 2012	Best Paper Award at EC'12
April 2010	Best Poster Award at WWW'10
August 2009–today	Fellowship by Swiss National Academic Foundation
May 2009	M.Sc. “with distinction” from KIT
November 2003–May 2008	Fellowship by German National Academic Foundation
June 2002	Award for excellent performance in school by the City of Aachen

Academic Visits and Internships

October 2012	**Microsoft Research NYC**, Host: Sebastien Lahaie
June 2012	**University of Cambridge**, Host: Felix Fischer
July 2011–September 2011	**Google Zürich**, Host: Radu Jurca
June 2010–August 2010	**Harvard University**, Host: David Parkes

Workshops and Doctoral Schools

April 2013	**Computation, Game Theory, and Economics**, Dagstuhl, Germany
Aug 2012	**Algorithmic Economics**, Pittsburgh, USA
September 2011	**New Trends in Mechanism Design**, Copenhagen, Denmark
January 2011	**Algorithmic Game Theory**, Los Angeles, USA
April 2010	**Computational Social Choice**, Estoril, Portugal

Peer-Reviewed Journal Publications

- **Payment Rules through Discriminant-Based Classifiers**
  Paul Dütting, Felix Fischer, Pichayut Jirapinyo, John Lai, Ben Lubin, and David Parkes
  Transactions on Economics and Computation (TEAC), invited to special issue
- **An Expressive Mechanism for Auctions on the Web**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  Transactions on Economics and Computation (TEAC), accepted subject to revision
- **Bidder Optimal Assignments for General Utilities**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  Theoretical Computer Science (TCS), Volume 478, Pages 22–32, March 2013
- **Sponsored Search, Market Equilibria, and the Hungarian Method**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  Information Processing Letters (IPL), Volume 113, Issue 3, Pages 67–73, February 2013
- **Offline File Assignments and Online Load Balancing**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  Information Processing Letters (IPL), Volume 111, Issue 4, Pages 178–183, January 2011

Peer-Reviewed Conference Publications

- **Auctions with Heterogeneous Items and Budget Limits**
  Paul Dütting, Monika Henzinger, and Martin Starnberger
  Workshop on Internet and Network Economics (WINE'12), Liverpool, UK, December 2012.
- **Maximizing Revenue from Strategic Recommendations under Decaying Trust**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  Conference on Information and Knowledge Management (CIKM'12), Maui, USA, December 2012
- **Payment Rules through Discriminant-Based Classifiers (Best Paper Award)**
  Paul Dütting, Felix Fischer, Pichayut Jirapinyo, John Lai, Ben Lubin, and David Parkes
  Conference on Electronic Commerce (EC'12), Valencia, Spain, June 2012
- **Simplicity-Expressiveness Tradeoffs in Mechanism Design**
  Paul Dütting, Felix Fischer, and David Parkes
  Conference on Electronic Commerce (EC'11), San Jose, USA, June 2011
- **An Expressive Mechanism for Auctions on the Web**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  World Wide Web Conference (WWW'11), Hyderabad, India, April 2011
- **How Much is Your Personal Recommendation Worth? (Best Poster Award)**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  World Wide Web Conference (WWW'10), Raleigh, USA, April 2010
- **Sponsored Search, Market Equilibria, and the Hungarian Method**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  Symposium on Theoretical Aspects of Computer Science (STACS'10), Nancy, France, March 2010
- **Bidder Optimal Assignments for General Utilities.**
  Paul Dütting, Monika Henzinger, and Ingmar Weber
  Workshop on Internet & Network Economics (WINE'09), Rome, Italy, December 2009

Invited Conference Publications

- **Mechanisms for the Marriage and Assignment Game**
  Paul Dütting and Monika Henzinger
  Conference on Algorithms and Complexity (CIAC'10), Rome, Italy, May 2010

Patent Applications

- **Network-Based Spam Detection**
  Paul Dütting and Radu Jurca
  Google, Zürich, Switzerland, December 2011

Academic Services

- **PC member:** IJCAI'13, EC'13, MATCH-UP '12

- **Reviewing for journals:** Algorithms, Information and Computation
- **Reviewing for conferences:** STOC'13, WINE '12, EC '12, EC '11, AAAI '11, ESA '06

Teaching Activities

- Algorithms (in English), Omid Etesami, EPFL, Winter 2012–2013
- Informatique II (in French), Jean-Cédric Chapellier, EPFL, Spring 2012
- Informatique I (in French), Ronan Boulic, EPFL, Spring 2011
- Informatik IV (in German), Hartmut Prautzsch, KIT, Spring 2007
- Informatik IV (in German), Rüdiger Dillmann, KIT, Spring 2006

References

Monika Henzinger

University of Vienna	Phone:	+43 1 427778310
Währinger Straße 29/6.32	E-mail:	monika.henzinger@univie.ac.at
A - 1090 Wien	Web:	http://www.cs.univie.ac.at/monika.henzinger/

Boi Faltings

EPFL IC LIA	Phone:	+41 21 6932735
INR 230, Station 14	E-mail:	boi.faltings@epfl.ch
CH - 1015 Lausanne	Web:	http://liawww.epfl.ch/~faltings/

David C. Parkes

Harvard University	Phone:	+1 617 3848130
33 Oxford Street	E-mail:	parkes@eecs.harvard.edu
Cambridge, MA 02138	Web:	http://www.eecs.harvard.edu/~parkes/