

Core-Pricing in Combinatorial Exchanges with Financially Constrained Buyers: Computational Hardness and Algorithmic Solutions

Martin Bichler, Stefan Waldherr

Department of Informatics, Technical University of Munich, Germany

The computation of market equilibria is a fundamental and practically relevant problem. Current advances in computational optimization allow for the organization of large combinatorial markets in the field. While we know the computational complexity and the types of price functions necessary for combinatorial exchanges with quasilinear preferences, the respective literature does not consider financially constrained buyers. We show that computing market outcomes that respect budget constraints, but are core-stable, is Σ_2^P -hard. Problems in this complexity class are rare, but ignoring budget constraints can lead to significant efficiency losses and instability as we demonstrate in this paper. We introduce mixed integer bilevel linear programs (MIBLP) to compute core-stable market outcomes, and provide effective column and constraint generation algorithms to solve these problems. While full core stability quickly becomes intractable, we show that realistic problem sizes can actually be solved if the designer limits attention to deviations of small coalitions. This n -coalition stability is a practical approach to tame the computational complexity of the general problem, and at the same time provides a reasonable level of stability for markets in the field where buyers have budget constraints.

Key words: multi-object auctions, combinatorial exchange, payment rules

1. Introduction

The analysis of market equilibria is arguably one of the most fundamental problems in the economic sciences. In the textbook model of perfect competition, a competitive equilibrium occurs when demand equals supply (Mas-Colell et al. 1995). The resulting price is often called the competitive price or market clearing price that will not change unless demand or supply changes. Participants have no incentive to change their behavior and the outcome is considered stable.

Most of the more recent micro-economic literature on competitive equilibria assumes a *utilitarian* or Benthamite welfare function, which maximizes the sum of all participants' utilities. This literature assumes an economy with purely quasilinear utility functions (i.e., bidders maximize payoff) and no budget constraints. However, ignoring budget constraints can lead to significant welfare losses and instability.

We draw on the core as the most prominent notion of stability. We aim to find core-stable outcomes in combinatorial markets with indivisible goods and general preferences that maximize welfare subject to budget constraints. Our main theoretical result proves that it is Σ_2^P -hard to find stable outcomes in combinatorial exchanges with payoff-maximizing but financially constrained participants. Then, we introduce algorithms to solve these problems. Despite the high computational complexity, we provide empirical evidence that stable solutions can be found for small but realistic problem sizes.

1.1. Combinatorial Markets

We focus on combinatorial markets, which allow bidders to specify package bids. This means a price is defined for a subset of the items. The bid price specified is only valid for the entire package and the package is indivisible such that bidders can express quasilinear preferences for general valuations including complements and substitutes. Combinatorial markets have found widespread application for the sale of spectrum licenses (Bichler and Goeree 2017), in truck-load transportation (Caplice and Sheffi 2006), for airport time-slots (Pellegrini et al. 2012, Ball et al. 2017), in day-ahead energy markets (Martin et al. 2014), for supply chain coordination (Fan et al. 2003, Guo et al. 2012, Walsh et al. 2000), and in transportation and industrial procurement (Schwind et al. 2009, Sandholm 2012). Milgrom (2007) has highlighted the importance of such markets for theory and practice.

Although package bids in a fully enumerative (XOR) bid language allow for general valuations, this generality comes at a price. First, the winner determination problem with an XOR (or also an OR) bid language becomes an *NP*-hard optimization problem (Pekec and Rothkopf 2003). Second, competitive equilibrium prices need to be non-linear and personalized to allow for maximum welfare (Bikhchandani and Ostroy 2002, Bichler and Waldherr 2017). Our definition of combinatorial markets in this paper is broad and includes all types of bid-languages that make the allocation problem a combinatorial optimization problem (Goetzendorff et al. 2015).

1.2. The Core

In an economy with purely quasilinear utilities and no budget constraints, a competitive equilibrium is defined as a feasible allocation and set of prices where buyers and sellers maximize payoff and the market is budget balanced (Bikhchandani and Ostroy 2002). In these economies the core coincides with the set of competitive equilibria, in which each participant maximizes payoff at the prices. The core is the set of feasible allocations that cannot be improved upon by a subset or coalition of the economy's agents if the coalition was to trade among each other. A core allocation consists of an assignment of items and its supporting price system and is weakly Pareto efficient. Weak Pareto efficiency describes an outcome for which there are no possible alternative outcomes whose realization would cause every participant to gain.

EXAMPLE 1. Consider an auction with a single seller s that sells a good without a reserve price. Suppose there are two buyers, b_1 with a value of \$5 for the good, and b_2 with a value of \$3 for the good. The core allocations consist of assigning the good to b_1 for any price p with $\$3 \leq p \leq \5 which results in a payoff of p for s , a payoff of $5 - p$ for b_1 and of 0 for b_2 . In these cases, no subset of the economy's agents can improve upon the outcome:

- i) The two buyers b_1, b_2 cannot both gain since they cannot trade with each other (hence both payoffs would be 0).
- ii) Trading with b_2 will lead to a payoff of at most \$3 for seller s , so s and b_2 cannot both gain from forming a coalition.
- iii) If s and b_1 would agree on a different price $p' \neq p$ for trading the good, it would either lower the payoff of s or b_1 , hence both cannot improve.

The core is the most prominent notion of stability and it can also be computed for markets with multiple objects for sale and multiple bidders (Day and Cramton 2012). The principle of core-stability is central for the computation of payments in high-stakes spectrum auction markets (Bichler and Goeree 2017). However, this literature assumes quasilinear utility functions, and bidders do not have budget constraints.

1.3. Exogenous Budget Constraints

Budget constraints are an important concern in most markets including spectrum auctions (Janssen et al. 2017), display ad auctions (Wu et al. 2018), and sponsored search auctions (Colini-Baldeschi et al. 2011). In most markets, bidders can only submit their budget-capped valuations, which can lead to significant inefficiencies as we show in this paper. Although the consideration of budget constraints appears as a practically important extension of the established quasilinear utility model, it turns out that it leads to substantial problems. Competitive equilibria in which each participant maximizes her payoff at the given prices might not be possible with budget constraints, as the package that maximizes the payoff of a buyer could not be affordable to her. It is also straightforward to see that a core-stable solution does not need to maximize welfare when bidders have budget constraints, see the following Example 2.

EXAMPLE 2. Suppose there are two buyers, b_1 and b_2 , having a value of \$10 and \$9 for a good, respectively. In addition, buyer b_1 has a budget constraint of \$1 and cannot spend more money. There are sellers s_1 and s_2 with reserve prices of \$0 and \$4. The welfare-maximizing allocation is to match b_1 and s_1 at a price between \$0 and \$1, and b_2 and s_2 at a price somewhere between \$9 and \$4, which yields \$15 gains from trade. However, this efficient allocation is not stable, because s_1 could approach b_2 and they could agree to deviate at a price of less than \$4 and more than \$1, which is profitable for both of them. Matching buyer s_1 to b_2 is stable, but does not allow b_1 to

trade with s_2 . Therefore, the gains from trade are only \$9, as compared to the welfare-maximizing allocation with gains from trade of \$15.

The example highlights an important impossibility: While core outcomes maximize welfare in Arrow-Debreu model (Arrow and Debreu 1954) and in markets with quasilinear utility but no budget constraints (Bikhchandani and Ostroy 2002), this no longer holds with budget constraints. Welfare-maximizing but unstable outcomes are hard to justify and maintain, which is why core constraints are enforced in high-stakes spectrum auctions. Therefore, we aim for *core-stable outcomes that maximize welfare subject to budget constraints*.

1.4. Contributions

Our paper has three main contributions. The first contribution is theoretical, as we provide a thorough analysis of the computational complexity of computing welfare-maximizing and core-stable outcomes of a combinatorial exchange with payoff-maximizing but financially constrained participants. This is a new and fundamental problem. We focus on combinatorial markets as they do not restrict the types of bidder valuations, and for their practical relevance. We prove that in the presence of payoff-maximizing, but budget-constrained buyers, the allocation and pricing problem becomes a Σ_2^p -hard optimization problem. This is important to show formally, but requires an elaborate reduction from the canonical Σ_2^p -hard problem QSAT₂. Problems in this complexity class are rare and considered intractable for all but toy problems. The hardness of these problems comes from the fact that prices and budgets need to be considered in the allocation problem, whereas with quasilinear preferences one can first solve the allocation and then the pricing problem (e.g. via core constraint generation as in Day and Raghavan (2007)). In addition, we also show restricted cases with dyadic coalitions that can be modeled and solved as integer programs and therefore are not Σ_2^p -hard, but in *NP*. Finally, we introduce n -coalition stability, describing a solution that is robust against coalitions of size at most n . This leads to a more tractable notion of stability as we will show. The cost of forming large coalitions with a dozen participants in a combinatorial market can be considered prohibitive in most markets. Note that it is already *NP*-hard for a given coalition with a given set of prices to determine whether a profitable deviation exists.¹

Second, we provide a quite general mixed integer bilevel linear program (MIBLP) to model the allocation and pricing problem of a combinatorial market with budget and core constraints, which can easily be adapted to specific bid languages.² While bilevel programming has been a topic in the literature for many years, algorithms to solve MIBLPs have only seen progress recently. Based on

¹ One can fix the allocation and set prices to zero. Then, finding a deviating coalition reduces to the standard winner determination problem in combinatorial auctions (Lehmann et al. 2006).

² Note that if we assume divisible goods then the model reduces from a MIBLP to a continuous bilevel program, which is known to be *NP*-hard. We do not discuss markets with divisible objects in this paper.

the MIBLP formulation, we develop column and constraint generation algorithms for combinatorial exchanges with budget constraints.

It is not obvious that realistic problem sizes of a Σ_2^P -hard problem could be solved in practice. In our third contribution, we perform extensive experimental analyses in which we provide evidence that even full core solutions can be found for small but realistic problem sizes. If we limit ourselves to n -coalition stability, we can compute much larger problem sizes while still providing a good level of stability for practical applications. We analyze two types of combinatorial markets, an airport time-slot allocation problem based on the CATS instance generator (Leyton-Brown et al. 2000), and a fishery rights exchange (Bichler et al. 2019). While full core stability quickly becomes intractable, we show that small but realistic problem sizes can actually be solved if the designer limits attention to deviations of coalitions with limited size. Our experimental results also show that if budget constraints are ignored and bidders can only submit budget-capped valuations, the computed prices and allocations are not stable and the welfare loss is substantial.

2. Related Work

Early on in the economic sciences, general equilibrium theory attempted to explain competitive equilibria in a market with multiple commodities. The Arrow–Debreu model shows that under convex preferences and perfect competition there must be a set of competitive equilibrium prices (Arrow and Debreu 1954). Market participants are price-takers, and they sell or buy goods in order to maximize their total value subject to their budget or initial wealth. The results derived from the Arrow–Debreu model led to the well-known welfare theorems, an important argument for markets as an efficient way to allocate resources. The first theorem states any competitive equilibrium leads to a Pareto efficient allocation of resources. The second theorem states that any efficient allocation can be attained by a competitive equilibrium, given the market mechanisms leading to redistribution.

In the Arrow–Debreu model each participant has an endowment of goods and money. Fisher markets are a simpler version of the Arrow–Debreu model, in which the total quantity of each product is given, and each buyer comes only with a monetary budget. They follow a tradition where utility is cardinal and individuals have interpersonally commensurable utility functions. These markets have received significant attention in the past twenty years in computer science, when researchers have gone beyond existence theorems and tried to find algorithms to actually compute allocations and prices (Vazirani 2007, Vazirani and Yannakakis 2011, Cole et al. 2017). In the Eisenberg–Gale convex program to solve Fisher markets, buyers have linear valuations that they aim to maximize subject to a budget constraint (Vazirani 2007). The designer maximizes a

Nash social welfare function, which is described as the budget-weighted geometric mean of the bidders' utilities.³

Most of the more recent micro-economic literature on competitive equilibria assumes a *utilitarian* or Benthamite welfare function, which maximizes the sum of all participants' utilities.⁴ As in the earlier literature on general equilibrium, the central question is when competitive equilibria exist. A number of authors explore conditions for linear and anonymous competitive equilibrium prices, so-called Walrasian equilibria (Kelso and Crawford 1982, Gul and Stacchetti 1999, Ausubel 2006, Leme 2017, Bichler et al. 2020). Bikhchandani and Mamer (1997) have shown necessary conditions for the aggregate valuation function of all individuals and Baldwin and Klemperer (2019) have characterized necessary conditions for the individual valuation functions to yield Walrasian equilibria. Baldwin et al. (2020) recently extended these results to markets with income effects, e.g. those with financial constraints. In addition, Bikhchandani and Ostroy (2002) discuss existence of non-linear and personalized competitive equilibrium prices. As indicated earlier, this entire literature assumes purely quasilinear utility functions and no budget constraints: market participants maximize their payoff, i.e. the value of an allocation minus the total price they pay for it. We continue to use these standard market design assumptions, but consider exogenous budget constraints that buyers might have.

Roughgarden and Talgam-Cohen (2015) highlight the tight connection between pricing, algorithms, and optimization and our work contributes to this line of research. We know that simple quasilinear preference models with unit demand allow for efficient computation and Walrasian prices, while more complex quasilinear preferences including complements and substitutes require the solution of NP-hard optimization problems and they demand non-linear and personalized prices. Table 1 relates different types of preferences to the type of price function necessary and the computational complexity to solve the allocation and pricing problem. In this paper, we show that the consideration of exogenous budget constraints yields an allocation and pricing problem that is even higher in the polynomial hierarchy compared to combinatorial exchanges with quasilinear utilities.

³ Often the literature assumes divisible goods. The Eisenberg-Gale program has been extended to accommodate indivisible objects (Cole et al. 2017) or separable, piecewise-linear concave utilities (Anari et al. 2018). With indivisible objects, the Nash social welfare maximization problem is NP-hard and APX-hard in general (Lee 2017), which has led to work on efficient approximation algorithms (Cole and Gkatzelis 2015, 2018). So far, however, the literature is restricted to relatively simple valuation functions.

⁴ The utilitarian welfare function is not only assumed in this theoretical literature, but also in spectrum auction markets (Bichler and Goeree 2016), in electricity markets (Madani and Van Vyve 2015), in markets for natural resource rights (Bichler et al. 2019), and most other market designs in the field. In a market with multiple buyers and sellers the prices cancel and the utilitarian welfare *maximizes the gains from trade*. We will also talk about *maximum welfare* or *maximum efficiency* in this case.

Table 1 Complexity Results and Price Functions for Computing a Core-Stable and Welfare Maximizing Outcome based on Different Types of Preferences

Preferences	Prices	Complexity	References
Unit demand	linear and anonymous	P	Shapley and Shubik (1971)
Strong substitutes	linear and anonymous	P	Milgrom and Strulovici (2009)
General quasilinear	non-linear and personalized	NP	Bikhchandani and Ostroy (2002)
General w. budgets	non-linear and personalized	Σ_2^P	<i>this paper</i>

Even though budget constraints have not been considered in the more recent competitive equilibrium theory, they have been a concern in other streams of the literature. Auction theory focuses on smaller auction markets with strategic bidders able to influence the price. Here, auctions are modeled as Bayesian games. While quasilinearity is also a standard assumption in this literature (Krishna 2009), a number of papers have dealt with the impact of budget constraints in auctions (Benoit and Krishna 2001, Borgs et al. 2005, Dütting et al. 2016). Unfortunately, it was shown that we cannot hope for any incentive-compatible mechanism in the presence of private budget constraints in multi-object auctions (Dobzinski et al. 2008, Colini-Baldeschi et al. 2011, Dütting et al. 2016), and there is a long literature addressing payoff-maximizing but budget-constrained bidders from a mechanism design perspective (Che and Gale 1998, Benoit and Krishna 2001, Pai and Vohra 2014). Many authors also deal with budget constraints in the context of dynamic advertising auctions (Conitzer et al. 2017, Borgs et al. 2007, Conitzer et al. 2018). Advertisers often have a budget for a campaign and they want to maximize payoff but consider their budget in a sequence of auctions. In contrast, we analyze a static environment with multiple buyers and sellers as is standard in competitive equilibrium theory, and want to compute prices that constitute a stable outcome where no coalition of budget-constrained buyers and sellers can deviate profitably.

3. Model and Preliminaries

In the following, we provide a formal definition of our model and concepts. We first introduce a model without budget constraints based on Bikhchandani and Ostroy (2002) and Bichler and Waldherr (2017). The papers show equivalence of the core and the set of competitive equilibria in a combinatorial exchange by drawing on specific linear programming formulations. The model without budget constraints is a convenient starting point for the analysis of budget constrained buyers.

There is a finite set of bidders N , consisting of buyers $i \in I$ and sellers $j \in J$ with $I \cup J = N$ and $I \cap J = \emptyset$, as well as a finite set of indivisible objects or items, K . Each buyer $i \in I$ has a non-negative value for each set of objects $S \subseteq K$ denoted $v_i(S) \in \mathbb{R}_{\geq 0}$ with $v_i(\emptyset) = 0$.

Sellers also have values or reservation prices for packages $Z \subseteq K$ with $v_i(Z) \in \mathbb{R}_{\geq 0}$. Buyers and sellers have free disposal. Every package is priced and each buyer $i \in I$ pays the price $p_i(S)$ for

the bundle S he receives, and each seller $j \in J$ receives the payment $p_j(Z)$ for the bundle Z he supplies. The vectors $P_i = (p_i(S))_{i,S}$ and $P_j = (p_j(Z))_{j,Z}$ describe the non-linear prices of buyers and sellers. In our initial analysis the preferences are quasilinear, i.e., the payoff of the buyer is $\pi_i = v_i(S) - p_i(S)$ and that of the seller is $\pi_j = p_j(Z) - v_j(Z)$. Later we will add budget constraints.

The problem of finding an efficient assignment maximizing gains from trade among buyers and sellers can be formulated as a linear program as follows: We use binary variables $x_i(S)$ to describe whether package S is assigned to bidder i and $y_j(Z)$ to describe whether package Z is supplied by seller j . The vectors $X = (x_i(S))_{i,S}$ and $Y = (y_j(Z))_{j,Z}$ describe the allocations of buyers and sellers. The model enumerates all possible allocations similar to the single-seller model in de Vries et al. (2007). The set of all possible object assignments is denoted as Γ , a specific assignment as $(X, Y) \in \Gamma$. For each possible allocation, we have a binary variable $\delta_{X,Y}$, which is 1 if an allocation is selected and 0 otherwise. The model allows for a very natural interpretation of the dual variable as prices. The dual variables of \mathbf{P} are written in brackets.

$$\begin{aligned}
w_P = \max \quad & \sum_{i \in I} \sum_{S \subseteq K} v_i(S) x_i(S) - \sum_{j \in J} \sum_{Z \subseteq K} v_j(Z) y_j(Z) \\
\text{s.t.} \quad & x_i(S) - \sum_{X: x_i(S)=1} \delta_{X,Y} = 0 && \forall i \in I, \forall S \subseteq K && (p_i(S)) \\
& -y_j(Z) + \sum_{Y: y_j(Z)=1} \delta_{X,Y} = 0 && \forall j \in J, \forall Z \subseteq K && (p_j(Z)) \\
& \sum_{S \subseteq K} x_i(S) \leq 1 && \forall i \in I && (\pi_i) \\
& \sum_{Z \subseteq K} y_j(Z) \leq 1 && \forall j \in J && (\pi_j) \\
& \sum_{(X,Y) \in \Gamma} \delta_{X,Y} = 1 && && (\pi_a) \\
& 0 \leq x_i(S) && \forall S \subseteq K, \forall i \in I \\
& 0 \leq y_j(Z) && \forall S \subseteq K, \forall j \in J \\
& 0 \leq \delta_{X,Y} && \forall (X, Y) \in \Gamma
\end{aligned} \tag{P}$$

The formulation \mathbf{P} introduces a variable $\delta_{(X,Y)}$ for each possible allocation, making the linear program large but integral. An LP solver is guaranteed to select a vertex such that $\delta_{(X,Y)} = 1$, such that we always get integer allocations $x_i(S)$ and $y_j(Z)$ of \mathbf{P} (Bichler and Waldherr 2017). At least one of these allocations maximizes the gains from trade, i.e., welfare in the economy. Note that even though \mathbf{P} is a linear program, its size is exponential in the number of bids since the linear program introduces a variable for each possible allocation. The underlying allocation problem in combinatorial exchanges is known to be NP-hard (Sandholm et al. 2002, Pekec and Rothkopf 2003). Even with the restriction of allowing only a single seller, the problem is equivalent to the winner determination problem in combinatorial auctions, which is also known to be NP-hard (Lehmann et al. 2006).⁵

Note that there are more effective formulations as binary programs that we will use in Section 5, where we discuss a bilevel program to compute core payments in the presence of financially

⁵ Note that even though the problem is NP-hard, there are algorithms that run in polynomial time in the size of the input if the number of bids is very large compared to the number of items (Lehmann et al. 2006).

constrained bidders. However, model **P** nicely shows how core payments can be computed without budget constraints.

The core-prices resulting from the dual variables of **P** are non-linear and personalized. We can now formulate the dual **D** of **P**.

$$\begin{aligned}
\min \quad & \sum_{i \in I} \pi_i + \sum_{j \in J} \pi_j + \pi_a \\
\text{s.t.} \quad & \pi_i \geq v_i(S) - p_i(S) & \forall i \in I, \forall S \subseteq K \quad (x_i(S)) \\
& \pi_j \geq p_j(Z) - v_j(Z) & \forall j \in J, \forall Z \subseteq K \quad (y_j(S)) \\
& \sum_{y_j(Z) \in Y} p_j(Z) - \sum_{x_i(S) \in X} p_i(S) + \pi_a \geq 0 & \forall (X, Y) \in \Gamma \quad (\delta_{X,Y}) \\
& \pi_i, \pi_j, p_i(S), p_j(Z) \geq 0 & \forall S, Z \subseteq K, \\
& & \forall i \in I, \forall j \in J \\
& \pi_a \in \mathbb{R}
\end{aligned} \tag{D}$$

The LP relaxation of the primal **P** is always integral such that strong duality holds. Now, the dual **D** introduces a price for each bidder and package, i.e., a non-linear and personalized price.

DEFINITION 1. Let $\Pi_i = (\pi_i) \in \mathbb{R}_{\geq 0}^{|I|}$ and $\Pi_j = (\pi_j) \in \mathbb{R}_{\geq 0}^{|J|}$ be the payoff vectors of the buyers and sellers in the auction \mathcal{E} and $V(\cdot)$ be the coalitional value function, i.e., the maximum transferable utility that can be gained by a coalition. Then (Π_i, Π_j) is in the *core* of the auction game \mathcal{E} , denoted $(\Pi_i, \Pi_j) \in \text{core}(\mathcal{E})$, if

$$\begin{aligned}
\sum_{i \in I} \pi_i + \sum_{j \in J} \pi_j &= V(N) && \text{core efficiency} \\
\sum_{i \in C} \pi_i + \sum_{j \in C} \pi_j &\geq V(C) \quad \forall C \subset N = I \cup J && \text{core rationality}
\end{aligned}$$

Bichler and Waldherr (2017) show that if $\pi_a = 0$, an optimal solution of **D** lies in the *core* of the auction, and the core is non-empty.

Quasilinear utility functions describe a game with transferable utility. The presence of budget constraints B_i of the buyers $i \in I$ violates quasilinearity, however, and only parts of the utility of the buyer up to the budget are transferable. We analyze the impact of this change in the following sections.

4. Complexity Analysis

In what follows, we analyze the complexity of welfare maximization subject to budget and core constraints. Example 2 has already illustrated that budget constraints can reduce the gains from trade. Even more issues can come up when bidders are not allowed to communicate their budget constraints, as is the standard in combinatorial auctions where bidders only submit their (non-restricted) bids.

EXAMPLE 3. Consider an example with two sellers, s_1 offering item A and s_2 offering item B and two buyers and their respective values and budgets as depicted in Table 2. Each buyer wants to obtain exactly one of the two items. If buyer b_1 is not allowed to communicate her budget and values, she can only place a bid of at most 4 for any of the items, leading to an assignment of

A to b_1 and B to b_2 which maximizes welfare for these reported values. However, this allocation is not stable since b_1 would approach s_2 to obtain item B in order to increase her true utility. Communicating financial constraints is necessary to ensure stability in this case.

Table 2 Example of values and budgets. If b_1 is not allowed to communicate her budget, the welfare-maximal allocation is unstable with regards to the true values.

	{A}	{B}	Budget
Buyer b_1	4	10	4
Buyer b_2	2	3	3

From Bikhchandani and Ostroy (2002) we know that the core of a combinatorial exchange can be empty even without budget constraints. But even if the core is non-empty without budget constraints, it can be empty if such constraints are added.

PROPOSITION 1. *A budget-constrained combinatorial exchange instance may have an empty core, even if the core is non-empty when budgets are ignored.*

Proof: Consider a case with two sellers, s_1 offering item A and s_2 offering item B and two buyers with values and budgets as demonstrated in Table 3. Without budget constraints, buyer b_1 can pay a price of 4 for each of the items, resulting in a core outcome. However, if we consider the budget constraints, there is no core outcome. Suppose buyer b_1 obtains $\{A, B\}$ for a combined price of at most 3. Then there is at least one seller with a payoff lower than 2, and buyer b_2 and this seller can form a coalition in which both are better off. Similarly, suppose b_2 obtains one of her desired items from one of the sellers while b_1 does not obtain any items. Since the combined payoff of the sellers is at most 2, b_1 and the sellers can form a coalition where all three are better off. Q.E.D.

Table 3 Example of values and budgets. There exists a stable allocation when no bidder is financially constrained, but the core is empty in the presence of budgets.

	{A}	{B}	{A,B}	Budget
Buyer b_1	0	0	10	3
Buyer b_2	4	4	4	2

The previous examples show that ignoring budget constraints can lead to substantial problems in combinatorial markets. Providing values and budget constraints in a market is not unusual. For example, in Google's auction for TV ads buyers were allowed to specify both (Nisan et al. 2009). However, allowing for budgets to be communicated comes at a price as well. In the following, we will show that it can be quite challenging to find core-stable outcomes in the presence of budget constraints.

A combinatorial exchange with budget constraints on the buyers' side can be seen as a game with partially transferable utility. The problem has a specific structure and therefore it is important to understand its computational complexity. We show that the problem is actually Σ_2^p -hard, a complexity class in the polynomial hierarchy that is higher than the class of NP-hard problems (Stockmeyer 1976).

THEOREM 1. *Computing a welfare-maximizing core outcome or providing a certificate that the core is empty in a combinatorial exchange with budget constraints is Σ_2^p -hard.*

Proof techniques for this complexity class are much less developed than those for lower levels in the polynomial hierarchy. The proof (see Appendix) reduces from QSAT_2 and requires an elaborate construction. A reduction from more abstract problems such as min max clique, which are known to be Σ_2^p -hard, appears simpler at first sight, but is not practical upon application.

The hardness of the problem comes from the fact that the allocation and pricing problem needs to be treated in a single optimization problem. Without budget constraints, the auctioneer can compute the welfare-maximizing solution first and then compute core-stable payments, as it is currently done for spectrum auctions (Day and Raghavan 2007). In this case, the welfare-maximal solution does always allow for prices such that the outcome is in the core (for example, charging the total bid amount). Thus, both calculations can be made independently. With budget constraints, however, this is not the case as there might not exist prices such that the welfare-maximizing solution can be extended to a core outcome.

An interesting question is whether there are some sufficient conditions for a combinatorial exchange to have a non-empty core that are simple to check. For transferable utility games the Bondareva-Shapley theorem describes balancedness as a *necessary and sufficient* condition for the core of a cooperative game with transferable utility (TU) to be non-empty (see Bondareva (1963) and Shapley (1967)). Balancedness is a rather obscure property, but it can be checked with linear programming in transferable utility games, and linear programming is also used to decide whether the core is empty in combinatorial exchanges without budget constraints (Bichler and Waldherr 2017).

A combinatorial exchange with budget constraints is closer to a game with non-transferable utility (NTU). The main result for NTU-games is that balanced games have a non-empty core, but the converse is not true (Scarf 1967). So, balancedness is a requirement that is sufficient for TU and NTU games to have a non-empty core. Scarf's algorithm is a central result to compute whether the core of an NTU game is empty. The computational version of Scarf's lemma is PPAD-complete (Kintali 2008). Moreover, the algorithm requires a matrix as an input that has a column with the payoff of each player for every coalition. In a combinatorial exchange with budget constraints,

every coalition can have multiple allocations with different payoffs. Moreover, we have partially transferable utility (up to the budget constraint), and the payoff vectors for each coalition are not unique for a coalition as is the case for a pure NTU game, but depend on the prices. In such price-guided markets Scarf’s algorithm does not provide a solution.

Therefore, it is interesting to understand restricted cases for which the problem falls into a lower complexity class. This approach has shown to be successful for standard multi-object markets with quasilinear bidders. For example, it is well-known that the gross substitutes condition is equivalent to $M^\#$ -concavity, a form of discrete concavity of a valuation function, which is a sufficient condition allowing for Walrasian prices and polynomial time algorithms to solve the allocation problem (Danilov et al. 2001, Milgrom and Strulovici 2009, Baldwin and Klemperer 2019).

One important case that simplifies the problem, however, is that of markets where we only care about dyadic coalitions, i.e., coalitions of size two, for which we introduce an integer program to solve in Section 6.3. Let us first introduce a bilevel optimization problem that allows us to solve the general case.

5. Optimization Model

Mixed integer bilevel linear programs (MIBLP) provide an adequate mathematical abstraction to model the allocation and pricing problem with budget-constrained bidders. Integer bilevel programs (IBLPs) are Σ_2^P -complete (Jeroslow 1985) as is our specific problem. Only recent algorithmic advances suggest that such problems can be solved in practice (Zeng and An 2014, Fischetti et al. 2017, Tahernejad et al. 2017).

MIBLPs belong to the class of bilevel optimization problems that have roots in the seminal work by Von Stackelberg (1934). Bilevel linear programs (BLPs) are frequently used to model sequential distributed decision-making. In these situations, typically a *leader* makes the first decision and a *follower* reacts after observing the leader’s decision. The follower’s action is important to the leader as it might interfere with the leader’s objective. The challenge of the leader is to predict the follower’s reaction and take action in such a way that after the follower’s reaction the leader’s objective is reached to the highest possible degree.

More technically, a BLP is a linear program that is constrained by another linear optimization problem. Usually the first optimization problem is called the upper level problem (leader) while the constraining problem is referred to as the lower level problem (follower). Given an upper level solution, the lower level computes an optimal solution under consideration of its respective constraints. This in turn affects the upper level by altering the value of the objective function or violating constraints, possibly making the overall solution infeasible. For this introduction to

bilevel optimization, let X be the set of variables in the upper level problem and Y be the set of variables in the lower level problem. Then, the general form of the problem is

$$\max_{x \in X} F(x, y) \tag{1a}$$

$$\text{s.t. } G(x, y) \leq 0 \tag{1b}$$

$$\min_{y \in Y} f(x, y) \tag{1c}$$

$$\text{s.t. } g(x, y) \leq 0 \tag{1d}$$

where $F, f: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^1, G: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^p, g: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^q$ are continuous, twice differentiable functions. Note that in MIBLP, F and f are represented by linear objective functions of the upper and lower level, while G and g are the respective linear constraints. BLPs, where X and Y include only continuous variables, are already *NP*-hard (Dempe 2002). A standard way to solve BLPs is to add the Karush-Kuhn-Tucker conditions of the lower-level program to the upper-level program. This adds complementarity constraints that can be modeled via integer variables. MIBLPs may include integer variables in the upper as well as in the lower level of the bilevel programming problem. In this case, the MIBLP cannot simply be formulated by modeling the KKT conditions of the lower level and the problems become very hard to solve.

The MIBLP we suggest finds core allocations in combinatorial exchanges with budget constraints. If the core is not empty, the solution of the bilevel program consists of a core allocation with maximum welfare. Additionally, we obtain prices and payments for buyers and sellers. If the core is empty, the MIBLP is infeasible. The goal of the upper level is to find prices $P = \{P_i, P_j\}$, such that the corresponding allocation (X, Y) is in the core and there is no core allocation with a higher welfare.

In the lower level of the bilevel program, we look for a coalition of bidders and sellers which can all improve their payoffs when trading among each other. For this, we introduce a variable $d \in \mathbb{R}$, indicating the *minimum improvement* a member of the coalition can achieve. Lower level variables $\chi_i(S) \in \{0, 1\}$ and $\gamma_j(Z) \in \{0, 1\}$ denote the allocation of bundles $S, Z \subseteq \mathcal{K}$ which are traded, and $\rho_i(S), \rho_j(Z)$ describe the corresponding prices and payments. Then, for each bidder that trades at least one item in the lower level, her difference in payoff (based on her trade in the upper and lower level) is calculated and considered for the minimum improvement d .

In the following, **(CE_x)** presents the general frame of the bilevel program that finds welfare-maximal core allocations and prices. For the moment, we ignore domain-specific allocation and pricing constraints. Hence, we simply denote the set of feasible allocations and prices by \mathcal{A} and \mathcal{P} in the upper level, and by \mathcal{A}_L and \mathcal{P}_L in the lower level, respectively. Depending on the exchange, these can represent arbitrary linear constraints that constitute a feasible trade. We assume an

XOR bidding language because it does not pose restrictions on the valuations. This means that \mathcal{A} contains a constraint that allows each bidder to only buy or sell a single bundle, respectively. This allows for a simpler representation of the welfare function as well as the calculation of the lower level deviation that should only depend on bidders actually participating in the blocking coalition. The general approach is not limited to XOR bidding languages, however, and can easily be adapted to other bid languages or exchanges that allow participants to buy and sell at the same time. We will give examples of how to adapt the general MIBLP (**CEx**) in Section 7. It is also straightforward to include swap bids that allow bidders to buy and sell items in a single bid.

$$\begin{aligned}
& \max_{x_i(S), y_j(Z)} \sum_{S \subseteq K} \sum_{i \in I} v_i(S) x_i(S) - \sum_{j \in J} \sum_{Z \subseteq K} v_j(Z) y_j(Z) && \text{(CEx)} \\
& \text{s.t. } \sum_{S \subseteq K} p_i(S) x_i(S) \leq \min \left\{ B_i, \sum_{S \subseteq K} v_i(S) x_i(S) \right\} && \forall i \in I \quad \text{(UBC)} \\
& p_j(Z) y_j(Z) \geq v_j(Z) y_j(Z) && \forall Z \subseteq K, \forall j \in J \quad \text{(UIRS)} \\
& x, y \in \mathcal{A} && \text{(UFA)} \\
& p \in \mathcal{P} && \text{(UFP)} \\
& d \leq 0 && \text{(Core)} \\
& d = \max d && \text{(Lower Level)} \\
& \text{s.t. } \sum_{S \subseteq K} \rho_i(S) \chi_i(S) \leq \min \left\{ B_i, \sum_{S \subseteq K} v_i(S) \chi_i(S) \right\} && \forall i \in I \quad \text{(LBC)} \\
& \rho_j(Z) \gamma_j(Z) \geq v_j(Z) \gamma_j(Z) && \forall j \in J, \forall Z \subseteq K \quad \text{(LIRS)} \\
& \chi, \gamma \in \mathcal{A}_L && \text{(LFA)} \\
& \rho \in \mathcal{P}_L && \text{(LFP)} \\
& d \leq \sum_{S \subseteq K} (v_i(S) - \rho_i(S)) \chi_i(S) - \\
& \quad \sum_{S \subseteq K} (v_i(S) - p_i(S)) x_i(S) + M \left(1 - \sum_{S \subseteq K} \chi_i(S) \right) && \forall i \in I \quad \text{(Imp-B)} \\
& d \leq \sum_{Z \subseteq K} \rho_j(Z) \gamma_j(Z) - \sum_{Z \subseteq K} p_j(Z) y_j(Z) + M \left(1 - \sum_{Z \subseteq K} \gamma_j(Z) \right) && \forall j \in J \quad \text{(Imp-S)} \\
& \sum_{i \in I} \sum_{S \subseteq K} \chi_i(S) + \sum_{j \in J} \sum_{Z \subseteq K} \gamma_j(Z) \geq 1 && \text{(Part)} \\
& \chi_i(S) \in \{0, 1\} && \forall S \subseteq K, i \in I \quad \text{(Binary)} \\
& \gamma_j(Z) \in \{0, 1\} && \forall Z \subseteq K, j \in J \quad \text{(Binary)} \\
& \rho_i(S) \in \mathbb{R}_0^+ && \forall S \subseteq K, i \in I \quad \text{(Real)} \\
& \rho_j(Z) \in \mathbb{R}_0^+ && \forall Z \subseteq K, j \in J \quad \text{(Real)} \\
& d \in \mathbb{R} && \text{(Real)}
\end{aligned}$$

$$\begin{array}{ll}
x_i(S) \in \{0, 1\} & \forall S \subseteq K, i \in I \quad (\text{Binary}) \\
y_j(Z) \in \{0, 1\} & \forall Z \subseteq K, j \in J \quad (\text{Binary}) \\
p_i(S) \in \mathbb{R}_0^+ & \forall S \subseteq K, i \in I \quad (\text{Real}) \\
p_j(Z) \in \mathbb{R}_0^+ & \forall Z \subseteq K, j \in J \quad (\text{Real})
\end{array}$$

The objective of **(CE \mathbf{x})** is to maximize gains from trade by determining an assignment of bundles and corresponding prices, such that the prices respect the budget constraints and individual rationality of buyers (BC) and sellers (UIRS). Furthermore, allocations and prices have to satisfy all domain-specific feasibility conditions for the allocation (UFA), and the prices (UFP). The prices have to be set in such a way that there is no coalition of participants which can benefit from deviating. To find such a coalition, an assignment χ, γ with payments ρ is determined in the lower level. Similar to the upper level, these assignments have to respect budget constraints (LBC), and individual rationality (LIRS) as well as all other feasibility conditions for assignments (LFA), and prices (LFP). Constraints (Imp-B) and (Imp-S) determine an upper bound for the minimum deviation of participants by calculating the improvement for each individual buyer and seller when participating in this coalition. Herein, M is a very big number to not restrict d in the case that a buyer or seller is not participating in a trade within the lower level (i.e., is not part of the blocking coalition). Constraint (Part) requires at least one participant in the lower level, a necessary condition since d would otherwise be trivially maximized by setting all χ, γ to zero in constraints (Imp-B) and (Imp-S). The objective of the lower level is to maximize the minimum improvement of participants of a blocking coalition. For the allocation (X, Y) and the corresponding payments to be in the core, this improvement must not be positive for any coalition (Core).

Note that the bi-linear term $\sum_{S \subseteq K} p_i(S)x_i(S)$ can easily be replaced by a single variable p_i (similar with p_j) for non-linear and personalized prices. In contrast, sometimes an auctioneer might want to have non-linear but anonymous prices and he could replace the variables $p_i(S)$ ($p_j(Z)$) for all i (j) by a single variable $p(S)$ ($p(Z)$) for each package S (Z). Neither personalized nor anonymous prices might be unique.

It is important to understand that the Σ_2^P -hardness of the problem hinges on the fact that items are indivisible. If the auctioneer had the opportunity to allocate fractional items or fractional packages, then both the lower- and the upper-level programs only contain rational allocation variables in $[0, 1]$ instead of binary variables, and the mixed-bilevel integer program reduces to a BLP.

Unfortunately, the core of the combinatorial exchange can also be empty, i.e., no allocation of items with prices for which there is no blocking coalition may exist. However, the gains for each blocking coalition might only be marginal and exceed the costs of finding such a blocking coalition

for the participants. In cooperative game theory, the ϵ -core is defined as the set of outcomes for which blocking coalitions can only improve by at most ϵ when deviating from the grand coalition. For the combinatorial exchange as defined above, the ϵ -core is equivalent to those allocations and prices for which there is no blocking coalition of buyers and sellers such that every member can improve its payoff by more than ϵ . A welfare-maximal ϵ -core outcome can be calculated by changing constraint (Core) to $d \leq \epsilon$.

6. Algorithms

If the lower level problem does not contain integer variables and is an LP, bilevel programs can be reformulated as single-level problem by replacing the lower level with its optimality conditions (e.g., Karush-Kuhn-Tucker (KKT)) and then solving the resulting mathematical program with equilibrium constraints, which can easily be translated into a mixed integer programming (MIP) problem and solved via standard MIP techniques (Bard and Moore 1990, Dempe 2002).

Bard and Moore (1990) initiated algorithmic solutions to mixed integer bilevel linear programs (MIBLPs). However, algorithms for MIBLPs are a relatively new research field, and there are no standard techniques for solving MIBLPs. MIBLPs with nonconvex lower level have been considered “still unsolved by the operations research community” (Delgado et al. 2010, Dempe 2003).

It has only been recently that two quite general branch-and-cut MIBLP algorithms have been proposed. Fischetti et al. (2017) propose an algorithm for MIBLPs with binary upper-level variables. They require the linking variables, those variables that have non-zero coefficients and are present in the upper- and lower-level program, to be integer. Tahernejad et al. (2017) propose another MIBLP solver based on cut generation which is available as open source in the MibS solver and has the same requirement on the linking variables.

A naïve technique would be to write down the full program with all optimality conditions of the lower level for all possible solutions of the upper level. However, this program would be huge and intractable for all but the smallest of toy instances. Even the smallest instances we consider in our computational experiments allow for 960 coalitions and hundreds of possible assignments for each of these coalitions. Hence, a naïve approach, enumerating all those options (and introducing KKT conditions for all of them) would be intractable to solve.

It is natural to use column and constraint generation techniques as the search space tends to be large. Zeng and An (2014) discuss a generic column and constraint generation framework. The authors first make use of the high-point relaxation of the bilevel program, wherein all lower level variables and constraints are duplicated into the upper level and a classical MILP is solved. This yields a solution that is feasible with respect to upper and lower level constraints, but not optimal with respect to the lower level. Actually, the high point relaxation is generally adopted as the

fundamental relaxation within MIBLP techniques in the literature (Moore and Bard 1990, Xu and Wang 2014, Scaparra and Church 2008). The solution serves as an upper bound UB for the optimal solution of the MIBLP.

A generic scheme for column and constraint generation in MIBLP, like that explored in Zeng and An (2014), can be summarized in the following steps:

1. Given an assignment x^* of the upper level variables in the single-level reformulation, the lower level problem is then solved to optimality, yielding an assignment y^* for the lower level variables.
2. If the combined solution (x^*, y^*) is feasible for the MIBLP, then $F(x^*, y^*)$ is a lower bound LB for its optimal solution. In the case that $LB = UB$, (x^*, y^*) is also an optimal solution.
3. Otherwise, let $y_{\mathbb{Z}} \in Y_{\mathbb{Z}}$ consist of the lower level variables with integer domain and $y_{\mathbb{R}} \in Y_{\mathbb{R}}$ denote the continuous lower level variables. The single-level reformulation is extended by KKT optimality conditions of the lower level with the integer variables fixed to $y_{\mathbb{Z}}^*$.

The procedure continues as described above, until lower bound and upper bound converge to the same value or the single-level reformulation is infeasible. Column and constraint generation is a wide-spread and textbook-level technique to solve large-scale integer programming problems. The challenge is how constraints and columns are generated and this is typically specific to the problem at hand.

In the case of computing welfare-maximal core allocations, the problem has a very special structure. For instance, simply adding the lower level variables and constraints to the upper level problem (i.e., using the high-point relaxation) does not restrict the upper level since the core feasibility constraint can be trivially satisfied (with $d = 0$) by setting the lower level allocation variables to the exact copy of the upper level allocation variables. Additionally, solving the lower level does not yield any meaningful lower bounds for the optimization problem: Either the lower level solution leads to an infeasibility, or the solution of the upper level is optimal.

In the following, we present an algorithm for computing welfare-maximal core allocations. As in the scheme by Zeng and An (2014) introduced above, our algorithm is also based on the classical idea of transforming the bilevel program into a single-level problem by dynamically adding columns and constraints, which are specific to our welfare maximization and pricing problem. The overall solution to the allocation and pricing problem involves the initial MIBLP formulation, an algorithm to solve the MIBLP formulation via column and constraint generation, the restriction on n -blocking-coalitions, and the delayed coalition generation discussed below. These components allowed us to solve realistic problem sizes as we show in our experiments.

6.1. Column and Constraint Generation

In what follows, we introduce our bilevel integer programming algorithm for the specific problem of computing a welfare-maximal core allocation and core prices in combinatorial exchanges.

First, the upper level is solved to optimality, ignoring the core constraint and possible blocking coalitions. Afterwards, given an optimal allocation x^*, y^* and prices p^* in the upper level, the lower level is solved to optimality for the corresponding optimal assignment of the upper level linking variables. If the optimal deviation d is at most zero, then there exists no coalition of bidders that block the upper level allocation and prices. In this case the upper level solution is feasible and, moreover, the welfare-maximal core allocation. If, however, the lower level yields a positive deviation d , let χ^*, γ^* be the allocation of items within the blocking coalition as determined by the lower level. Then, the upper level is extended by the Karush-Kuhn-Tucker (KKT) optimality conditions of the lower level with the integer variables (i.e., χ and γ) fixed to the result of the lower level. These additional constraints force the upper level to determine a new allocation and / or new prices such that there is no possibility for a blocking coalition with an allocation χ^*, γ^* to determine payments among each other such that all members of the coalition can profit.

Algorithm 1 illustrates the overall procedure. A welfare-maximal upper level solution (x, y, p) considering the current constraints is determined in Line 3 of the while loop. If the upper level is infeasible and no allocation has been fixed, then there exists no stable outcome. Otherwise, if the allocation in the upper level was fixed before, no prices for this allocation exist. Hence, the allocation is unfixed in Line 9 and this allocation and all other allocations with higher social welfare are forbidden since they were proven to be unstable in previous iterations. Afterwards, the upper level is solved again. If the upper level is feasible, we solve the lower level in order to find a deviating coalition that can improve upon the upper level allocation and prices. If there is no such coalition, (x, y, p) constitute welfare-maximal core allocations and prices. Otherwise, we add the KKT condition for the fixed lower level integer variables to the upper level in Line 17 and fix the current welfare-maximal allocation in Line 18.

Algorithm 1: Bilevel algorithm to obtain welfare-maximal core allocation and prices

```

1 Set fixedAllocation = false;
2 while true do
3   Solve the upper level  $U$  to obtain allocation  $x, y$  and prices  $p$ ;
4   if  $U$  is infeasible then
5     if fixedAllocation = false then
6       return that there exist no core allocation and prices;
7     else
8       Add a constraint to  $U$  to forbid allocation  $(x, y)$  as well as all allocations with
          higher social welfare;
9       Set fixedAllocation = false;
10      Continue;
11    end
12  end
13  Solve the lower level  $L$  with allocations  $x, y$  to obtain deviation  $d$  and allocation  $\chi, \gamma$ ;
14  if  $d < 0$  and fixedAllocation == false then
15    return  $x, y, p$  as the welfare-maximal core allocation and prices
16  else
17    Add KKT condition for allocation  $\chi, \gamma$  to  $U$  ;
18    Fix allocation  $x, y$  in  $U$  ;
19    Set fixedAllocation = true;
20  end
21 end

```

Algorithm 1 already offers a sufficient framework to obtain welfare-maximal core allocations and prices. However, the nature of the problem and its inherent complexity can lead to very long runtimes. In the following, we introduce several methods to further improve our algorithm in order to obtain solutions faster, or at all. First, in Section 6.2, we discuss some computational approaches such as branching schemes and reducing the size of the IPs that need to be solved in Algorithm 1. Afterwards, we address the computational challenges by restricting the number and size of blocking coalitions. As we show in Section 6.3, with blocking coalitions of size 2, the problem is even in NP (and hence no longer Σ_2^P -hard). Based on these ideas, we introduce delayed coalition generation in Section 6.5.

6.2. Computational Approaches

First, we discuss two speedup strategies to improve Algorithm 1. It should be noted that the effect of these improvement depends on the structure of specific instances and can lead to large improvements in some cases while having only little effect in others. Hence, while discussing these two general strategies, we also lay out in which cases these strategies lead to the biggest improvements.

Branching over possible payments for fixed allocations. In each level, branching can be defined by selecting a subset of bidders B and a payment threshold p' dividing the problem into two subproblems, one using the constraint $\sum_{b \in B} \sum_{S \subset K} p_b(S) \geq p'$ and one using the constraint $\sum_{b \in B} \sum_{S \subset K} p_b(S) \leq p'$. Since for a fixed allocation, the precise payments have no impact on the

welfare, we can terminate the algorithm as soon as we find a core-stable outcome for any of the subproblems. In instances where a welfare-maximal core allocation is easy to find (i.e., not many allocations have to be forbidden in Line 9), but finding prices is very difficult, this branching procedure led to noticeable speedups in determining whether an allocation can be supported by prices to result in a stable outcome. However, in cases where it is easy to find blocking coalitions for the allocations determined in the upper level regardless of prices, adding additional constraints via a branching scheme to differentiate between prices can even lead to unnecessary computational overhead.

Removal of old KKT conditions. Another modification to speed up the algorithm is by not only adding new constraints and variables in each iteration in the form of KKT conditions, but to also remove old ones regularly in order to keep the upper level problem from becoming too large. Deviating coalitions and the allocation of items among them depend significantly on the upper level solution. Thus, the KKT conditions that were added iteratively for a fixed allocation x, y in the upper level, become redundant when the allocation could be proven to not allow for stable prices and a new allocation x', y' is determined. In this case, the coalitions that blocked the previous outcome might no longer form based on the new allocation. Then, instead of keeping the corresponding KKT conditions for the remainder of the solution process, these constraints and variables can be removed and a singular constraint can be introduced which prohibits the allocation that was proven not to support core prices. This strategy is very effective when allocations with high welfare are very different from another (i.e., very different allocations are fixed in Line 18 in successive rounds of Algorithm 1). On the other hand, if these allocations are very similar, they also lead to similar blocking coalitions and hence a re-creation of the removed KKT conditions.

6.3. Exclusion of Dyadic Coalitions

A fundamental problem of the MIBLP is the exponential number of coalitions that could possibly block the upper level solutions. The concept of the core considers coalitions of any size. Large coalitions are costly, not only in terms of tractability of the bilevel program, but also for participants to find. The computation of the coalitional value of each possible coalition is *NP*-hard to compute in general.

An alternative approach is to focus on coalitions of restricted size. For some applications, it might be sufficient to find a solution that avoids deviations of dyadic coalitions.

For this, we introduce variables $\rho_{ij}(S) \in \{0, 1\}$ for each possible package trade between a buyer i and seller j over all packages $S \subset K(j)$ where $K(j)$ denotes all bundles $Z \subseteq K$ which are offered by j . The variable ρ is set to 1, whenever the respective dyad would form a blocking coalition.

Similar to the general problem, we introduce constraints such that only outcomes without blocking coalitions are feasible.

$$\begin{aligned}
& \max_{x_i(S), y_j(Z)} \sum_{S \subseteq K} \sum_{i \in I} v_i(S) x_i(S) - \sum_{j \in J} \sum_{Z \subseteq K} v_j(Z) y_j(Z) \rho_{ijS} & (\mathbf{DY}) \\
& \text{s.t. } \sum_{S \subseteq K} p_i(S) x_i(S) \leq \min \left\{ B_i, \sum_{S \subseteq K} v_i(S) x_i(S) \right\} & \forall i \in I \quad (\text{BC}) \\
& p_j(Z) y_j(Z) \geq v_j(Z) y_j(Z) & \forall Z \subseteq K, \forall j \in J \quad (\text{IRS}) \\
& \sum_{i \in I} \sum_{S \subseteq K} p_i(S) x_i(S) = \sum_{j \in J} \sum_{Z \subseteq K} p_j(Z) y_j(Z) & (\text{BB}) \\
& \sum_{S: k \in K} \sum_{i \in I} x_i(S) \leq \sum_{Z: k \in K} \sum_{j \in J} y_j(Z) & \forall k \in K \quad (\text{Supply}) \\
& \sum_{S \subseteq K} x_i(S) \leq 1 & \forall i \in I \quad (\text{XOR-B}) \\
& \sum_{Z \subseteq K} y_j(Z) \leq 1 & \forall j \in J \quad (\text{XOR-S}) \\
& \pi_i = \sum_{S \subseteq K} (v_i(S) - p_i(S)) x_i(S) & \forall i \in I \quad (\text{payoffB}) \\
& \pi_j = \sum_{Z \subseteq K} (p_j(Z) - v_j(Z)) y_j(Z) & \forall j \in J \quad (\text{payoffS}) \\
& \pi_j + v_j(S) \geq B_i \delta_{ij}(S) & \forall i \in I, \forall j \in J, \forall S \subseteq K(j) \quad (\text{Block-B}) \\
& \pi_j + v_j(S) \geq v_i(S) - \pi_i - M \gamma_{ij}(S) & \forall i \in I, \forall j \in J, \forall S \subseteq K(j) \quad (\text{Block-Imp}) \\
& \delta_{ij}(S) \geq \gamma_{ij}(S) & \forall i \in I, \forall j \in J, \forall S \subseteq K(j) \quad (\text{No-Block}) \\
& x_i(S) \in \{0, 1\} & \forall S \subseteq K, i \in I \quad (\text{Binary}) \\
& y_j(Z) \in \{0, 1\} & \forall Z \subseteq K, j \in J \quad (\text{Binary}) \\
& \delta_{ij}(S), \gamma_{ij}(S), \rho_{ij}(S) \in \{0, 1\} & \forall i \in I, \forall j \in J, \forall S \subseteq K(j) \quad (\text{Binary}) \\
& \pi_i, p_i(S) \in \mathbb{R}_0^+ & \forall S \subseteq K, i \in I \quad (\text{Real}) \\
& \pi_j, p_j(Z) \in \mathbb{R}_0^+ & \forall Z \subseteq K, j \in J \quad (\text{Real})
\end{aligned}$$

Constraints (Block-B) to (No-Block) of (\mathbf{DY}) characterize blocking dyads and require some explanation. Note that a buyer i wants to deviate if for her new payoff $v_i(S) - p_{ij}(S) > \pi_i$ would hold, where $p_{ij}(S)$ is some transfer price in a blocking dyad. Similarly, a seller j would want to deviate if $p_{ij}(S) - v_j(S) > \pi_j$. Rearranging terms, $\pi_j + v_j(S) < v_i(S) - \pi_i$ characterizes a blocking coalition, i.e., with $\pi_j + v_j(S) \geq v_i(S) - \pi_i$ a dyad would not be blocking (see (Block-Imp)). We also need to consider budget constraints of buyers B_i . With $\pi_j + v_j(S) \geq B_i$ in constraint (Block-B), we avoid payments to the seller j characterized by the LHS of the constraint that are higher than the budget of the buyer B_i . The binary variable $\gamma_{ij}(S) = 1$ indicates if a dyad would deviate

due to improvement in payoffs, variable $\delta_{ij}(S) = 1$ if the required payments would exceed budget. Constraint (No-Block) demands that a dyad can only be willing to deviate due to payoffs if the required payments would exceed the budget of the buyer involved, since this dyad would otherwise be blocking the outcome.

Note that without the constraints (BC), (Block-B), (Block-Imp), and (No-Block), we have the winner determination problem in a combinatorial exchange, which is known to be NP-hard. With these additional constraints involving additional binary variables, the problem cannot be solved in polynomial time. Interestingly, however, the problem is not Σ_2^P -hard anymore, but is in *NP* since it can be solved by the integer program mixed integer program (**DY**).

6.4. Restrictions on Blocking Coalitions

Beyond dyadic coalitions, one can restrict the cardinality of coalitions to those with only a few participants. We will refer to such outcomes where we only achieve stability against blocking coalitions with at most n bidders as *n-coalition stable* or in short *n-stable* outcomes. Such a restriction speeds up the solution process in two places. First, the solution space of the lower level program becomes smaller due to the additional constraints on coalitions and possible omission of participants. Second, the smaller number of prospective deviating coalitions obviously leads to fewer options for deviations and hence to fewer KKT conditions that need to be added to the upper level before obtaining welfare-optimal outcomes which are *n-coalition stable*. In a similar way, one can leverage prior information about likely coalitions and find outcomes that are stable with respect to these coalitions. In many cases, it might be sufficient to find *n-coalition stable* outcomes or outcomes that are stable against coalitions of close participants due to the high computational or organizational cost (or even inability) for these bidders to find blocking coalitions of larger sizes, themselves.

6.5. Delayed Coalition Generation

As discussed in the previous sections, stable outcomes against subsets of coalitions are easier to find than stable outcomes against all possible blocking coalitions. In delayed coalition generation (DCG), instead of considering all coalitions and their possible trades from the beginning, we determine an initial set of coalitions $\mathcal{C} \subseteq C$ and solve the MIBLP, only considering coalitions in \mathcal{C} . For example, we can search for *n-coalition stable* outcomes by defining \mathcal{C} as the set of all coalitions with size at most n . If we find an outcome that is stable against coalitions in \mathcal{C} , then we extend \mathcal{C} in order to obtain stability with regards to a larger set of coalitions. If, however, one cannot find a stable outcome against coalitions in \mathcal{C} , the auction cannot be stable in general. DCG can also be seen as a procedure to determine the largest size of blocking coalitions, for which a stable outcome can exist.

7. Experimental Design

At first sight the problem appears to be too hard to solve even small problems. Interestingly, we show that small but realistic problem sizes can actually be solved at the present time. Given the advances in mixed integer bilevel programming and hardware in the last ten years, we expect to solve increasingly larger problem sizes in the future.

In order to evaluate the empirical hardness of the problem, we draw on two problem types with different characteristics, a combinatorial exchange for the allocation of airport time slots to airlines, and a market for fishery access rights. For the airport time slot market, we could handle instances up to 8 airports, 40 airlines and 80 slots traded. In the fishery access rights market we could solve problems with 10 sellers and 10 buyers and hundreds of units traded of two distinct items (share classes). The bidding languages differ, but the MIBLP is very flexible and can easily be adapted on the upper and lower level to the very specifics of the allocation problem. In both environments, we restricted the analysis to a computation time of 10 minutes and for some instances to two hours. First, this allowed us to run a large number of experiments. Second, we found that even computation times of 24 hours typically did not get much better results.

7.1. Airport Time Slots

The first market we consider is the allocation of airport time slots which has been discussed repeatedly in the literature (Rassenti et al. 1982, Castelli et al. 2011, Pellegrini et al. 2012). Current assignment mechanisms for slots have come under much scrutiny over the years. Theoretically, airport time slots are reallocated at the start of each season. However, within the current slot allocation process, airlines in Europe and the US enjoy grandfather rights over the slots they obtained in previous seasons which leads to inefficient usage. For example, in 2016, only 22 slots were made available for auction by the Heathrow airport. Ball et al. (2017) argue that currently there is a strong case for the use of market mechanisms to allocate or reallocate such slots. Package bids are essential in this domain: A takeoff slot is only valuable with a landing slot. Such slots can cost millions of US dollars. Haylen and Butcher (2017) report airlines that bid up to 75 million US dollars for a pair of slots at the Heathrow airport. Airlines are interested in many of these combinations, but there are concerns about depressed bidding of smaller airlines that are financially constrained (Debyser 2016).

7.1.1. Data

For our experiments we draw on the valuation model for airport slot allocation in the combinatorial auction test suite (CATS) (Leyton-Brown et al. 2000), a very widely used instance generator for combinatorial auctions. Given the input of a number of airports, their respective location, and available slots, CATS generates values for pairs of slots at two distinct airports, representing a

flight between these two respective cities. Each bidder is interested in obtaining one pair of slots. The values for these pairs are based on a common value for each slot at a given airport as well as a private deviation for each bidder. The instances we generate for our market describe bids by up to 40 bidders that are interested in up to 80 slots at the eight coordinated airports. In addition to the values generated by CATS, we randomly generated budgets for buyers which lie between the buyer’s highest value and half of this amount. In the generated instances, the bidders truthfully bid these values and budgets. All problem instances are available online.⁶

7.1.2. Domain-Specific Adaptations for the MIBLP

In the airport slot allocation model as described above, buyers submit bids for various packages of two items where each of the items is owned by a seller (airport). Because each trade of slots requires at least one buyer and two sellers, it is not possible for dyadic coalitions to block an allocation. Each buyer is only interested in obtaining one of these packages, i.e., the buyers use an XOR bidding language. Hence, the sets \mathcal{A} and \mathcal{A}_L in the MIPBL of Section 5 describe all allocations where each buyer is assigned at most one pair of two slots. Sellers can freely dispose all slots that are not assigned to buyers.

7.2. Fishery Access Rights

The second market we consider is one for fishery access rights (catch shares) that was recently implemented in Australia (Bichler et al. 2019). A catch share describes the right to catch a certain volume of a specific type of fish in a specific region. After a reform, some fishers needed more catch shares while others wanted to sell some or even all their endowment and exit the market. This made package bidding a necessity, because sellers who wanted to exit did not want to sell only part of their shares. There were concerns about buyers being financially constrained such that they could not bid up to their net present value for shares. We analyze an exchange design where we do consider budget constraints explicitly and one where bidders can only submit budget-capped values and explore problem sizes we can solve.

In this exchange, there exists a set of share classes \mathcal{L} which are traded among a set I of buyers and J of sellers. Each buyer $i \in I$ can submit multiple bids, one for each share class and can win any combination of these bids (OR bids). A bid for share class $l \in \mathcal{L}$ by a buyer has to include a lower and an upper bound \underline{X}_{il} , and \bar{X}_{il} of the desired units of this share class and a value v_{il} for a single unit of l . However, each seller $j \in J$ submits a single bid for a bundle of share classes containing the number of units per share class in the bundle and an ask price. Sellers are only interested in selling their entire bundle.

⁶ <http://anonymized.com>

7.2.1. Data

We were fortunate to draw on an instance generator used to evaluate the scalability of a combinatorial exchange design that was later used in the field (Bichler et al. 2018). The instance generator is based on information of the real-world market with regards to the licenses each registered fishing business owns and the estimated revenue generated by these shares calculated by landings and fish market prices. Given this field data, a distribution of share classes and values was generated. For a specified number of bidders, sellers, and share classes, the generator then constructed corresponding instances based on these distributions. While these are synthetic instances smaller than those described in Bichler et al. (2018), care was taken to closely reflect the specifics of the market, the very bidders that participated, their endowments, and historical catch levels. The instance generator allowed us to simulate small and large exchanges by taking only subsets of the participants and share classes into account.

Based on historical market data, the generator simulates values for share classes and generates bids (i.e., the lower and upper bound $\underline{X}_{il}, \bar{X}_{il}$ of shares requested by the buyer, as described in the previous section). Budget constraints were not part of the generator, and hence we added them based on the following assumption. We assumed the upper bounds on requested shares simulated by the generator to be induced by an underlying financial constraint: A buyer would bid for a larger number of shares but could not afford to pay up to the reported valuation for them. Therefore, we set the budget at the amount that a buyer would pay if she was to obtain all of her requested shares. Afterward, we multiplied the upper bounds of all requested share classes by a random factor between 1.3 and 2.0. This led to the scenario outlined above: Buyers requested more shares, but were only able to pay up to their values for a smaller fraction of shares. Again, all instances used for our experiments are available online.

7.2.2. Domain-Specific Adaptations for the MIBLP

In order to adapt the MIBLP program of Section 5, for each buyer $i \in I$ the variables $x_i(S)$ can be replaced by variables $x_{il} \in \mathbb{N}$ where x_{il} is the number of units per share class $l \in L$ that buyer i receives. Then, terms $\sum_{S \subseteq K} v_i x_i(S)$ can be replaced by $\sum_{l \in \mathcal{L}} v_{il} x_{il}$. Likewise, for each seller $j \in J$, the variables $y_j(Z)$ can be replaced by y_j since they only offer a single bundle. The variables χ and γ of the lower level can be replaced in the same way. Further, the lower and upper bounds of share classes for each buyer have to be included in the constraint sets \mathcal{A} and \mathcal{A}_L and the constraint (Imp-B) in the lower level has to be adapted in order to implement an OR rather than an XOR bidding language for the buyers.

8. Results

The number of participants and items were the main treatment variables in both sets of experiments. Runtime for the core computations and also welfare gains compared to markets where buyers can only provide capped values were the main focus variables.

For our experiments we used a 24 core Intel Xeon ES-2620 (2.00 GHz) with 64 GB memory on Ubuntu 19.04.01, using Gurobi 8.1.0 for solving the mixed integer linear programs.

8.1. Airport Time Slots

Let us first discuss the results for the exchange of airport time slots for which we used CATS to generate several instances for the 8 coordinated airports (referred to as *sellers*). Treatment variables for our experiments include the number of airlines that submit bids (referred to as *buyers*), the number of slots that are sold by the airports (referred to as *items*), and the size n of blocking coalitions for which we can reach n -coalition stability. First, we present results for the performance of our algorithm with regard to runtime and size of instances that can still be solved. Then, we discuss the negative effects of not considering the budget constraints within the computation.

We consider the market for airport slot allocation with 8 sellers, between 40 and 80 items that were endowed uniformly across all sellers, and 10 to 50 buyers. These scenarios cover markets that range from little competition to those with a large amount of competition over the items and the number of possible blocking coalitions becomes quite large as shown in Table 4. Note that each blocking coalition must include at least two sellers (since buyers are only interested in bundles of items from two distinct sellers) and contain at most two sellers for each buyer (since buyers are only interested in a bundle).

#buyers	#items	number of coalitions of		
		size 3	size 5	unrestricted size
10	8	280	7,420	182,350
25	8	700	90,300	$\approx 8 \cdot 10^8$
40	8	1,120	343,280	$\approx 2 \cdot 10^{14}$
50	8	1,400	653,100	$\approx 2 \cdot 10^{17}$

Table 4 Number of coalitions of up to a specific size for the problem sets in the airport market.

8.1.1. n-Coalition Stability For each combination of buyers and number of items, we generated 8 instances and tried to determine welfare-maximal core allocations and prices with Algorithm 1 within a time limit of 10 minutes. We also applied the algorithm to find allocations and prices that are n -coalition stable for $n = \{3, 5\}$. In Table 5 we report for each treatment combination, for how many of these 8 instances Algorithm 1 returned a solution within that time limit (i.e., how

many of the instances were “solved”). Noticeably, for all instances in which the algorithm terminated, a stable outcome could be found. Further, we report the mean runtime to find a solution (in those cases where a solution was found) and its standard error.

#buyers	#items	3-stable			5-stable			core-stable		
		solved	time mean	std.err.	solved	time mean	std.err.	solved	time mean	std.err.
10	40	8	0.32	0.02	8	0.77	0.05	8	3.03	0.86
10	80	8	5.22	0.83	8	8.37	1.84	8	15.58	7.37
25	40	8	1.42	0.23	8	5.23	1.10	8	98.03	27.26
25	80	8	24.93	3.95	8	48.76	9.79	2	293.84	0.25
40	40	8	2.04	0.12	8	7.68	0.93	6	222.89	30.09
40	80	8	50.01	6.61	8	93.06	16.28	0	–	–
50	40	8	3.07	0.27	8	11.85	0.96	5	425.53	52.21
50	80	8	69.24	4.88	8	83.68	5.96	0	–	–

Table 5 Computational results of the airport slot market for a variable number of buyers and sellers. For 8 instances each, the number of instances for which outcomes were calculated that are n -coalition stable and the average computation time in seconds required to solve the instances is shown.

Even for larger sized problems with 50 buyers and 40 items, we were able to solve 5 out of 8 instances within the 10 minute time limit. It should be noted that in the 3 remaining cases, as well as 6 out of 8 cases for 40 buyers and 80 items, the algorithm was able to derive core-stable solutions when allowed a time limit of two hours. However, for the largest instances with 50 buyers and 80 items, even allowing for a two hour time limit did not yield core-stable outcomes. Such instances rarely find the optimal solution even if run overnight.

For all instances we were able to derive n -coalition stable outcomes for coalitions up to size 5 within at most two minutes of runtime. Despite the computational complexity of the problem, these results show that we are able to find core-stable outcomes in realistic markets.

In Table 6, we report results for delayed coalition generation (DCG). Since coalitions of size $n = 2$ can never block any assignment, we start with coalitions of size $n = 3$. As shown in Table 5, 3-coalition stable outcomes can be computed in very short time. Whenever an outcome is found that is n -coalition stable, we increase n and search for an $n + 1$ -coalition stable outcome (or prove that no such outcome exists). We have again allowed for a time limit of 10 minutes. For these experiments, we have concentrated on the harder cases with 40 and 50 buyers. Table 6 shows the number of instances for which we could prove whether the problem was core-stable within 10 minutes. If the algorithm could not prove core stability within this time, we report the maximal n for which n -coalition stability was proven within the 10 minutes averaged over the instances that could not be solved completely.

#buyers	#items	solved	avg. n -stability of not solved
40	40	5	23.67
40	80	0	8.62
50	40	6	23.00
50	80	0	8.13

Table 6 Results of delayed coalition generation for the airport slot market

Even in instances where no core-stable solution could be obtained without DCG, DCG could find outcomes that were n -coalition stable for large n . For all instances, DCG could assure 7-coalition stability and for 80 items an average n -coalitional stability of $n \geq 8$ could be achieved. For 50 buyers and 40 items, DCG was able to find one additional core-stable solution compared to the results of the experiments reported in Table 5.

In all other instances with 40 items full core stability could not be proven, but DCG resulted in an average n -coalition stability of $n \geq 23$. In most applications it will be very costly for coalitions of 23 participants to form, and robustness against coalitions of this size would be considered very stable.

Longer computation times did not lead to significantly larger problem instances that could be solved. However, with advances in computational optimization, we expect to see progress made on these problems in the future.

8.1.2. Welfare Gains An important question concerns the welfare gains one can expect when considering values and budgets rather than just values capped by the budget constraints of a bidder, or when ignoring budgets at all. In the first setting, which we refer to as *capped bidding*, bidders submit their values up to the budget limit only and the auctioneer computes bidder-optimal core-selecting prices based on these capped values. In a second *unrestricted* setting, bidders and the auctioneer ignore their budget constraints and bid up to their true valuations as if there were no budget constraints hoping that the prices are within budget.

Table 7 shows the negative effects of ignoring financial constraints. As compared to the resulting allocation and prices when reporting the budget constraints, in 21 of the 24 instances different allocations emerged. The gains from trade were on average 8.09 percent lower with capped bidding. When buyers submitted unrestricted bids up to their valuation, this resulted in violations of their budget constraints in 16 out of the 24 instances, i.e., bidders actually made a loss.

Overall, a simple bid language that does not let buyers express valuations and overall budgets can lead to significantly lower welfare and instability.

8.2. Fishery Access Rights

For experiments in the fishery market, treatment variables include the number of fishers that want to add additional shares (referred to as *buyers*), the number of fisheries that want to sell their shares

#buyers	#items	capped bidding			unrestricted bidding
		instances with different allocations	average welfare loss	std. error welfare loss	instances leading to losses
10	40	6	6.30 %	2.40 %	3
10	80	8	16.98 %	1.98 %	5
25	40	7	4.66 %	1.46 %	8
in total		21	9.31 %	1.58 %	16

Table 7 Negative effects of ignoring financial constraints: Welfare loss in case of capped bidding; instances with prices leading to losses in case of unrestricted bidding

(referred to as *sellers*), and the size n of blocking coalitions for which we can reach n -coalition stability. In contrast to the market for airport slots in the previous section, where airlines interested in operating flights between two airports needed two items (slots) from these two distinct airports, fishers who want to buy shares need a larger quantity of items (shares) of the same share class which they can purchase from an arbitrary number of sellers. We consider a fishery market with two share classes of protected fish, 5 to 10 buyers, and 5 to 15 sellers. The valuations of share classes for all participants as well as the number of shares in these two classes that were demanded by buyers and offered by sellers was determined by the generator based on real-world data described in Section 7.2.1. For the smaller instances with 5 buyers and 5 sellers, the total number of shares available in the market was between 209 and 430 while for the larger instances with 10 buyers and 15 sellers it was between 650 and 1050. The variance in the number of shares is due to the generator that created small but realistic fishery markets with significant competition among the participants. The number of possible blocking coalitions grows quickly as is shown in Table 8.

#buyers	#items	number of coalitions of		
		size 3	size 5	unrestricted size
5	5	125	575	960
5	10	375	4,275	31,712
5	15	750	16,725	1,015,776
10	10	1,000	20,425	1,046,528
10	15	1,875	62,825	33,520,640

Table 8 Number of coalitions of up to a specific size for the problem sets in the fishery market.

8.2.1. N-Coalition Stability As in the experiments for the airport slot market, we generated 8 instances for each combination of buyers and sellers and tried to determine welfare-maximal core allocations and prices with Algorithm 1 within a time limit of 10 minutes. We also applied the algorithm to find allocations and prices that are n -coalition stable for $n = \{3, 5\}$. In contrast to the airport slot market in Section 8.1, not all instances that we were able to solve allowed for core-stable outcomes. In Table 9 we report for each treatment combination, how many of these 8

instances could be solved within that time limit, and how many of these instances are n -coalition stable. Further, the average time to find a solution (in those cases where a solution was found) and the corresponding standard error are presented. For example, with 5 buyers and 10 sellers only 4 out of 7 problem instances that could be solved were actually 5-coalition stable, for the other 3 instances we could show that no such allocation and prices exist.

#buyers	#sellers	3-stable				5-stable				core-stable			
		solved	stable	time		solved	stable	time		solved	stable	time	
				mean	std.err.			mean	std.err.			mean	std.err.
5	5	8	8	0.35	0.10	8	8	1.73	0.51	8	8	2.16	0.73
5	10	8	8	1.40	0.55	7	4	57.07	24.21	6	3	25.07	14.18
5	15	8	8	18.73	10.7	7	7	11.43	6.41	7	6	14.33	3.98
10	5	8	8	13.63	8.24	5	5	39.39	18.83	5	5	51.12	21.53
10	10	6	6	25.48	12.31	3	3	278.00	63.02	1	1	321.59	–
10	15	3	3	367.02	61.60	0	0	–	–	0	0	–	–

Table 9 Computational results of the fishery market for a variable number of buyers and sellers. For 8 instances each, the number of instances for which outcomes were calculated that are n -coalition stable and the average computational time required to solve the instances in seconds is shown.

In Table 10, we report results for DCG. We start with coalitions of size $n = 2$, for which results can be obtained relatively easy compared to larger sizes of coalitions (cf. Section 6.3), and increase n whenever an outcome is found that is n -coalition stable. Again, we allowed for a time limit of 10 minutes. The table shows the number of instances for which we could prove within the time limit either (a) that the problem was core-stable (i.e., for which we found core-stable allocations and prices), or (b) that the problem did not allow for a core-stable allocation. In both cases, we considered the corresponding problem instance to be “solved” by our algorithm. We also describe for how many of the instances we were able to derive 3- and 5-coalition stable outcomes and full core stability (“core”).

Moreover, we analyze those instances which could not be solved by our algorithm (“not solved”). For these, in the last column of Table 10, we report the maximal n for which n -coalition stability was proven within the 10 minutes averaged over the instances that could not be solved completely. For example, in the case of 10 buyers and 10 sellers, none of the 8 instances could be solved, but we were able to show robustness against coalitions of size 3.5 on average.

With DCG we were able to provide a solution for one additional instance for 10 buyers and 5 sellers when compared to the experiments described in Table 9 (where we could prove that there are no stable outcomes against coalitions of up to 6 members). Also, we were able to find 3-coalition stable solutions when starting DCG with coalitions of size 2 as opposed to only considering coalitions of size 3.

#buyers	#sellers	solved	3-stable	5-stable	core	not solved	avg. n -stability of not solved
5	5	8	8	8	8	0	–
5	10	6	8	4	3	2	4
5	15	7	8	7	6	1	3
10	5	6	8	5	5	2	3
10	10	0	8	1	0	8	3.5
10	15	0	4	0	0	8	2.5

Table 10 Results of delayed coalition generation for the test set with two share classes and a variable number of buyers and sellers

In larger instances especially, it was harder to leverage DCG in order to prove n -coalition stability for higher n . For example, we were not able to find the core-stable solution for 10 buyers and 10 sellers that we could find when directly looking for outcomes that are stable against coalitions of arbitrary size. However, we were able to find 2-coalition stable allocations for all instances and on average 2.5-coalition stable outcomes for the largest instances with DCG. Again, longer computation times did not lead to significantly larger problem instances that could be solved.

8.2.2. Welfare Gains Like we did for the airport slot market, we also evaluated the negative effects of ignoring the budget constraints in the fishery market. Again, we considered capped bidding (i.e., bidders submit their budget capped valuations) and unrestricted bidding (i.e., bidders ignore their budget constraints and bid up to their valuations). For the test instances for which we were able to calculate 3-coalition stable outcomes in all instances, we compared these results to the two alternatives in which budgets cannot be communicated.

Table 11 summarizes the results. As compared to the 3-stable outcome, the gains from trade were on average 32.17% lower with capped bidding. Moreover, in 9 of the 32 instances different allocations emerged, not only with respect to the number of share classes traded per bidder, but with respect to which buyers traded which share classes. When buyers submitted unrestricted bids up to their valuation, this resulted in violations of their budget constraints in 17 out of the 32 instances, i.e., bidders actually made a loss.

#buyers	#items	capped bidding			unrestricted bidding
		instances with different allocations	average welfare loss	std. error welfare loss	instances leading to losses
5	5	2	35.48 %	3.62 %	5
5	10	1	37.33 %	2.20 %	5
5	15	0	36.54 %	2.29 %	0
10	5	6	17.52 %	3.70 %	7
in total		9	32.17 %	2.02 %	17

Table 11 Negative effects of ignoring financial constraints: Welfare loss in case of capped bidding; instances with prices leading to losses in case of unrestricted bidding

The market for fishery access rights shows that there can be substantial welfare losses when only capped valuations are taken into account and bidders cannot properly express valuations and prices. Capped valuations allow buyers to only bid for a lesser number of items in this market, while in the airport slot market package bids are only on pairs of slots. However, even in the airport time slot auction the welfare loss was significant.

9. Conclusions

In this paper, we have analyzed combinatorial exchanges in the presence of financially constrained bidders. We have shown that ignoring budgets and allowing bidders only to submit budget capped values leads to welfare losses and instability, even though the outcome was stable with respect to the capped values. In order to maximize welfare subject to stability of the market, a market designer therefore should take budget constraints into account.

We have proven that computing welfare maximizing and core-stable outcomes leads to a Σ_2^P -hard optimization problem. Problems in this level of the polynomial hierarchy are rare in business practice and are typically considered intractable. This is an important insight and contributes to the literature connecting algorithm design and market design. Although the computational hardness of the problem could be interpreted as an impossibility result at first sight, we were able to solve markets with restricted forms of core stability and smaller problem instances even to full core stability.

To this end, we have introduced a bilevel integer programming formulation and effective algorithms to compute welfare maximizing outcomes that are robust against coalitions of restricted size, i.e., n -coalition stability. While bilevel integer programs have been discussed for several decades, they have received more attention in the past five years. Still, there are no established black-box approaches as they are available for integer programming nowadays. In this paper, we have provided a column and constraint generation framework which is tailored to our problem and surprisingly effective in solving realistic problem instances. In particular, delayed column generation provided an effective way to systematically increase the size of the coalitions against which a computed outcome is stable. An outcome that maximizes welfare, considers budget constraints, and is robust against deviations of coalitions with four or five participants is computable even for large markets. This might well be sufficiently stable in practice, because it is computationally hard for larger coalitions to evaluate whether they can find a profitable deviation.

Given the advances in computational optimization in the past three decades we expect to solve increasingly larger problem sizes, such that we can hope to find core-stable solutions in real-world environments where budget constraints matter. In addition, it will be interesting to further explore specific types of restricted bidder preferences and allocation problems that have a lower computational complexity and allow for more efficient computation.

References


- Anari, Nima, Tung Mai, Shayan Oveis Gharan, Vijay V Vazirani. 2018. Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2274–2290.
- Arrow, K. J., G. Debreu. 1954. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society* 265–290.
- Ausubel, Lawrence M. 2006. An efficient dynamic auction for heterogeneous commodities. *American Economic Review* **96**(3) 602–629.
- Baldwin, Elizabeth, Omer Edhan, Ravi Jagadeesan, Paul Klemperer, Alexander Teytelboym. 2020. The equilibrium existence duality: Equilibrium with indivisibilities & income effects.
- Baldwin, Elizabeth, Paul Klemperer. 2019. Understanding preferences: ‘demand types’, and the existence of equilibrium with indivisibilities. *Econometrica* **to appear**.
- Ball, Michael O, Frank Berardino, Mark Hansen. 2017. The use of auctions for allocating airport access rights. *Transportation Research Part A: Policy and Practice* .
- Bard, Jonathan F, James T Moore. 1990. A branch and bound algorithm for the bilevel programming problem. *SIAM Journal on Scientific and Statistical Computing* **11**(2) 281–292.
- Benoit, Jean-Pierre, Vijay Krishna. 2001. Multiple-object auctions with budget constrained bidders. *The Review of Economic Studies* **68**(1) 155–179.
- Bichler, M., V. Fux, J. Goeree. 2018. A matter of equality: Linear pricing in combinatorial exchanges. *INFORMS Information Systems Research* **28**(to appear).
- Bichler, M., J. Goeree. 2016. Frontiers in spectrum auction design. *International Journal of Industrial Organization* **forthcoming**.
- Bichler, Martin, Maximilian Fichtl, Gregor Schwarz. 2020. Walrasian equilibria from an optimization perspective: A guide to the literature. *TUM Working Paper* **1**(1).
- Bichler, Martin, Vladimir Fux, Jacob K Goeree. 2019. Designing combinatorial exchanges for the reallocation of resource rights. *Proceedings of the National Academy of Sciences* **116**(3) 786–791.
- Bichler, Martin, J Goeree. 2017. *Handbook of Spectrum Auction Design*. Cambridge University Press.
- Bichler, Martin, Stefan Waldherr. 2017. Core and pricing equilibria in combinatorial exchanges. *Economics Letters* .
- Bikhchandani, Sushil, John W Mamer. 1997. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of economic theory* **74**(2) 385–413.
- Bikhchandani, Sushil, Joseph M Ostroy. 2002. The package assignment model. *Journal of Economic theory* **107**(2) 377–406.

- Bondareva, Olga N. 1963. Some applications of linear programming methods to the theory of cooperative games. *Problemy kibernetiki* **10** 119–139.
- Borgs, Christian, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, Mohammad Mahdian. 2007. Dynamics of bid optimization in online advertisement auctions. *Proceedings of the 16th international conference on World Wide Web*. ACM, 531–540.
- Borgs, Christian, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, Amin Saberi. 2005. Multi-unit auctions with budget-constrained bidders. *Proceedings of the 6th ACM conference on Electronic commerce*. ACM, 44–51.
- Caplice, Chris, Yossi Sheffi. 2006. Combinatorial auctions for truckload transportation. *Combinatorial auctions* **21** 539–571.
- Castelli, Lorenzo, Paola Pellegrini, Raffaele Pesenti. 2011. Airport slot allocation in europe: economic efficiency and fairness. *International journal of revenue management* **6**(1-2) 28–44.
- Che, Yeon-Koo, Ian Gale. 1998. Standard auctions with financially constrained bidders. *The Review of Economic Studies* **65**(1) 1–21.
- Cole, Richard, Nikhil Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay V Vazirani, Sadra Yazdanbod. 2017. Convex program duality, fisher markets, and nash social welfare. *Proceedings of the 2017 ACM Conference on Economics and Computation*. 459–460.
- Cole, Richard, Vasilis Gkatzelis. 2015. Approximating the nash social welfare with indivisible items. *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. 371–380.
- Cole, Richard, Vasilis Gkatzelis. 2018. Approximating the nash social welfare with indivisible items. *SIAM Journal on Computing* **47**(3) 1211–1236.
- Colini-Baldeschi, Riccardo, Monika Henzinger, Stefano Leonardi, Martin Starnberger. 2011. On multiple round sponsored search auctions with budgets. *arXiv preprint arXiv:1112.6361* .
- Conitzer, Vincent, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolas E Stier-Moses, Chris Wilkens. 2018. Pacing equilibrium in first-price auction markets. *arXiv preprint arXiv:1811.07166* .
- Conitzer, Vincent, Christian Kroer, Eric Sodomka, Nicolas E Stier-Moses. 2017. Multiplicative pacing equilibria in auction markets. *arXiv preprint arXiv:1706.07151* .
- Danilov, Vladimir, Gleb Koshevoy, Kazuo Murota. 2001. Discrete convexity and equilibria in economies with indivisible goods and money. *Mathematical Social Sciences* **41**(3) 251–273.
- Day, R., P. Cramton. 2012. The quadratic core-selecting payment rule for combinatorial auctions. *Operations Research* **60** 588–603.
- Day, R., S. Raghavan. 2007. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science* **53** 1389–1406.

- de Vries, Sven, James Schummer, Rakesh V Vohra. 2007. On ascending vickrey auctions for heterogeneous objects. *Journal of Economic Theory* **132**(1) 95–118.
- Debysier, Ariane. 2016. Airports in the EU: Challenges Ahead. Tech. rep., European Parliamentary Research Service. doi:10.2861/149777.
- Delgadillo, Andrés, José Manuel Arroyo, Natalia Alguacil. 2010. Analysis of electric grid interdiction with line switching. *IEEE Transactions on Power Systems* **25**(2) 633–641.
- Dempe, Stephan. 2002. *Foundations of bilevel programming*. Springer Science & Business Media.
- Dempe, Stephan. 2003. Annotated bibliography on bilevel programming and mathematical programs with equilibrium constraints .
- Dobzinski, Shahar, Ron Lavi, Noam Nisan. 2008. Multi-unit auctions with budget limits. *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*. IEEE, 260–269.
- Dütting, Paul, Monika Henzinger, Martin Starnberger. 2016. Auctions for heterogeneous items and budget limits. *ACM Transactions on Economics and Computation* **4**(1) 4.
- Fan, Ming, Jan Stallaert, Andrew B Whinston. 2003. Decentralized mechanism design for supply chain organizations using an auction market. *Information Systems Research* **14**(1) 1–22.
- Fischetti, Matteo, Ivana Ljubić, Michele Monaci, Markus Sinnl. 2017. A new general-purpose algorithm for mixed-integer bilevel linear programs. *Operations Research* .
- Goetzendorff, A., M. Bichler, P. Shabalín, R. W. Day. 2015. Compact bid languages and core pricing in large multi-item auctions. *Management Science* **61**(7) 1684–1703.
- Gul, F., E. Stacchetti. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* **87** 95–124.
- Guo, Zhiling, Gary J Koehler, Andrew B Whinston. 2012. A computational analysis of bundle trading markets design for distributed resource allocation. *Information Systems Research* **23**(3-part-1) 823–843.
- Haylen, Andrew, Louise Butcher. 2017. Airport slots. Tech. rep., UK House of Commons Briefing Paper No CBP 488.
- Janssen, Maarten, Vladimir A. Karamychev, Bernhard Kasberger. 2017. *Budget Constraints in Combinatorial Clock Auctions*. Cambridge University Press, 318–337.
- Jeroslow, Robert G. 1985. The polynomial hierarchy and a simple model for competitive analysis. *Mathematical programming* **32**(2) 146–164.
- Kelso, A. S., V. P. Crawford. 1982. Job matching, coalition formation , and gross substitute. *Econometrica* **50** 1483–1504.
- Kintali, Shiva. 2008. Complexity of scarf’s lemma and related problems. *arXiv preprint arXiv:0812.1601* .
- Krishna, Vijay. 2009. *Auction theory*. Academic press.

- Lee, Euiwoong. 2017. Apx-hardness of maximizing nash social welfare with indivisible items. *Information Processing Letters* **122** 17–20.
- Lehmann, Daniel, Rudolf Müller, Tuomas Sandholm. 2006. The winner determination problem. *Combinatorial auctions* 297–318.
- Leme, Renato Paes. 2017. Gross substitutability: An algorithmic survey. *Games and Economic Behavior* **106** 294–316.
- Leyton-Brown, Kevin, Mark Pearson, Yoav Shoham. 2000. Towards a universal test suite for combinatorial auction algorithms. *Proceedings of the 2nd ACM conference on Electronic commerce*. ACM, 66–76.
- Madani, Mehdi, Mathieu Van Vyve. 2015. Computationally efficient mip formulation and algorithms for european day-ahead electricity market auctions. *European Journal of Operational Research* **242**(2) 580–593.
- Martin, Alexander, Johannes C Müller, Sebastian Pokutta. 2014. Strict linear prices in non-convex european day-ahead electricity markets. *Optimization Methods and Software* **29**(1) 189–221.
- Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green, et al. 1995. *Microeconomic theory*, vol. 1. Oxford university press New York.
- Milgrom, Paul. 2007. Package auctions and exchanges. *Econometrica* **75**(4) 935–965.
- Milgrom, Paul, Bruno Strulovici. 2009. Substitute goods, auctions, and equilibrium. *Journal of Economic theory* **144**(1) 212–247.
- Moore, James T, Jonathan F Bard. 1990. The mixed integer linear bilevel programming problem. *Operations research* **38**(5) 911–921.
- Nisan, Noam, Jason Bayer, Deepak Chandra, Tal Franji, Robert Gardner, Yossi Matias, Neil Rhodes, Misha Seltzer, Danny Tom, Hal Varian, et al. 2009. Google’s auction for TV ads. *International Colloquium on Automata, Languages, and Programming*. Springer, 309–327.
- Pai, Mallesh M, Rakesh Vohra. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* **150** 383–425.
- Pekec, A., M. H. Rothkopf. 2003. Combinatorial auction design. *Management Science* **49**(11) 1485–1503.
- Pellegrini, Paola, Lorenzo Castelli, Raffaele Pesenti. 2012. Secondary trading of airport slots as a combinatorial exchange. *Transportation Research Part E: Logistics and Transportation Review* **48**(5) 1009–1022.
- Rassenti, Stephen J, Vernon L Smith, Robert L Bulfin. 1982. A combinatorial auction mechanism for airport time slot allocation. *The Bell Journal of Economics* 402–417.
- Roughgarden, Tim, Inbal Talgam-Cohen. 2015. Why prices need algorithms. *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 19–36.
- Sandholm, T. 2012. Very-large-scale generalized combinatorial multi-attribute auctions: Lessons from conducting \$60 billion of sourcing. *the handbook of market design*.

-
- Sandholm, Tuomas, Subhash Suri, Andrew Gilpin, David Levine. 2002. Winner determination in combinatorial auction generalizations. *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*. ACM, 69–76.
- Scaparra, Maria P, Richard L Church. 2008. A bilevel mixed-integer program for critical infrastructure protection planning. *Computers & Operations Research* **35**(6) 1905–1923.
- Scarf, Herbert E. 1967. The core of an n person game. *Econometrica: Journal of the Econometric Society* 50–69.
- Schwind, Michael, Oleg Gujo, Jens Vykoukal. 2009. A combinatorial intra-enterprise exchange for logistics services. *Information systems and e-business management* **7**(4) 447–471.
- Shapley, L. S, M. Shubik. 1971. The assignment game I: The core. *International Journal of game theory* **1**(1) 111–130.
- Shapley, Lloyd S. 1967. On balanced sets and cores. *Naval Research Logistics (NRL)* **14**(4) 453–460.
- Stockmeyer, Larry J. 1976. The polynomial-time hierarchy. *Theoretical Computer Science* **3**(1) 1–22.
- Tahernejad, Sahar, Ted K Ralphs, Scott T DeNegre. 2017. A branch-and-cut algorithm for mixed integer bilevel linear optimization problems and its implementation. *Working Paper* .
- Vazirani, Vijay V. 2007. Combinatorial algorithms for market equilibria. *Algorithmic Game Theory* 103–134.
- Vazirani, Vijay V, Mihalis Yannakakis. 2011. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM (JACM)* **58**(3) 10.
- Von Stackelberg, Heinrich. 1934. *Marktform und Gleichgewicht*. Springer.
- Walsh, W. E., M. P. Wellman, F. Ygge. 2000. Combinatorial auctions for supply chain formation. *Proceedings of the 2Nd ACM Conference on Electronic Commerce*. EC '00, ACM, New York, NY, USA, 260–269. doi:10.1145/352871.352900. URL <http://doi.acm.org/10.1145/352871.352900>.
- Wu, Di, Xiujun Chen, Xun Yang, Hao Wang, Qing Tan, Xiaoxun Zhang, Jian Xu, Kun Gai. 2018. Budget constrained bidding by model-free reinforcement learning in display advertising. *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. 1443–1451.
- Xu, Pan, Lizhi Wang. 2014. An exact algorithm for the bilevel mixed integer linear programming problem under three simplifying assumptions. *Computers & operations research* **41** 309–318.
- Zeng, Bo, Yu An. 2014. Solving bilevel mixed integer program by reformulations and decomposition. *Optimization online* 1–34.

Figure 1 Illustration of buyers' interests, only concerning items of type ψ 

Appendix A: Complexity Analysis

In the following we prove that finding a welfare-maximizing core allocation with exogenous budget constraints is Σ_2^p -complete by a reduction from the canonical Σ_2^p -complete problem QSAT₂.

2-Quantified Satisfiability, QSAT₂: Given a $n + m$ variable Boolean formula $\varphi(x, y)$ in DNF with $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_m)$ is it true that $\exists x \forall y \varphi(x, y)$?

A.1. Membership in Σ_2^p

We first prove that the problem of finding a core outcome of welfare D is in the class Σ_2^p . Let $x(S), y(Z), p(S), p(Z)$ be a certificate for the allocations and prices. The gains from trades can be easily verified in polynomial time by using this certificate. Further, showing that this outcome is in the core is in *co-NP* since any blocking coalition C with corresponding assignments $\chi^C(S), \gamma^C(Z), p^C(S), p^C(Z)$ is a certificate that the outcome is not in the core.

A.2. Idea behind the transformation

Before formally proving the theorem, we give a short explanation of the reduction and the ensuing relationship between an instance of QSAT₂ and the corresponding combinatorial exchange. We concentrate on the main items and buyers with a direct correspondence to the world of QSAT₂ and omit the various auxiliary items, buyers and sellers. For these, we refer to the complete description of the transformation below.

In the combinatorial exchange, we define items relating to the truth assignment of x and y variables as well as the truth values which clauses evaluate to. For the variables x and y , items χ and γ are introduced and the truth assignment of variables x and y in QSAT₂ depends on which of the buyers obtains these items. Each clause is represented by n^2 items of type ψ which will indicate whether the clause evaluates to true or false, again depending on which buyers obtain which of these items.

We introduce different types of buyers. For $i \leq n$, buyers B_i^K and B_i^M , each concerned with the items corresponding to the truth assignments of variable x_i and the clauses affected by it. For $j \leq m$, buyers B_j^G , which are concerned with the items corresponding to the truth assignments of y_j and the clauses affected by it. The construction is such that either all buyers of type B^K win one of their preferred packages or they do not win any items. In the former case, the corresponding instance of QSAT₂ evaluates to true, in the latter case it is false.

The connection between χ, γ and ψ variables in the exchange and the correspondence of setting clauses to false by assigning truth values to variables in QSAT₂ is done via defining bundles of items in which the buyers are interested in. Figure 1 demonstrates the situation, showing the items ψ corresponding to three clauses in form of a matrix (we will refer to these as clause matrices in the following). Additionally, items χ, γ and bundles in which buyers of the various types are interested in, are shown. Buyer B_i^K is interested in either χ_i or $\bar{\chi}_i$ items as well as the i -th ‘row’ of one clause matrix. More formally, he is interested in the bundle

$$(\chi_i \vee \bar{\chi}_i) \wedge (\{\psi_{1i1}, \dots, \psi_{1in}\} \vee \{\psi_{2i1}, \dots, \psi_{2in}\} \vee \dots \vee \{\psi_{Li1}, \dots, \psi_{Lin}\}),$$

where L is the number of clauses. Buyers B_i^M are interested in buying one out of χ_i or $\bar{\chi}_i$ as well as the i -th ‘column’ in all clause matrices of clauses which include the corresponding x_i or \bar{x}_i variable. In Figure 1, a bundle for buyer B_2^M , including χ_2 and the second column of the first clause matrix (since C_1 is the only clause containing x_2) is depicted. Finally, buyers B_j^G are interested in bundles which contain one item out of γ_j or $\bar{\gamma}_j$ and complete clause matrices for clauses which include the corresponding y_j or \bar{y}_j variables. As can be seen, the individual bundles block each other and can not be obtained simultaneously for each clause matrix. The corresponding clause evaluates to true if and only if neither buyers of type B^M buys a column of the matrix or buyers of type B^G buy the complete matrix. For example, in Figure 1, no items of the second clause matrix are won by either a buyer type B^M or B^G . In this case, buyers of type B^K can all obtain their respective row of the (second) clause matrix. Consequently, for this example, the second clause and therefore the entire expression evaluates to true.

The valuations and budgets of buyers are defined in such a way that buyers of type B^K have the highest value for their respective bundles, but only small budgets which does not allow them to bid up to their true valuation. In contrast, buyers B^M have high valuations and sufficient budget to buy the bundle they are interested in. Buyers B^G have low valuations and can not compete with buyers B^M . However, their budget is high enough in order to outbid buyers B^K . In order to obtain sufficiently high welfare gains, buyers B^K must obtain their desired bundles (i.e. win one of the clause matrices) and the outcome must be stable such that B^G and the sellers do not want to deviate by assigning the items to buyers B^G or B^M instead. Each buyer of type B^K can only obtain one of his desired bundles containing at least one row in one clause matrix (see Figure 1, which is equivalent to the corresponding clause evaluating to true in QSAT₂), when no other buyer purchases a column within this matrix.

Buyers B^K and B^M are designed in such a way that B_i^K obtains the χ_i -item corresponding to the truth assignment of x_i and B_i^M its negation. Thus, buyers B^M obtain the columns in each clause matrix relating to the clauses which are set to false due to the truth assignment of variables x . Because of their lower valuations and budgets, buyers B^G can only compete for columns in clause matrices corresponding to clauses not yet set to false due to the assignment of x . These buyers maximize their payoffs when they can purchase as many complete matrices as possible which are not blocked by buyers B^M . In QSAT₂ this corresponds to assigning truth values to variables y in such a way that as many as possible of the remaining clauses evaluate to false (i.e. those which are not already evaluating to false due to the assignment of x variables). Only if the buyers of type B^G cannot manage to block all remaining clause matrices (the y variables in QSAT₂),

buyers B^K can purchase rows in at least one of the matrices (the x variables in QSAT_2) relating to one clause which evaluates to true. Then, the assignments of items corresponding to truth values of x is a solution for the QSAT_2 problem. In other words, if B^K win in every allocation, then there exists a stable outcome that achieves the pre-defined welfare in the decision problem.

A.3. Transformation

We present a transformation with valuations using an XOR bidding language. The transformation can easily be done for an OR bidding language as well, however this requires additional auxiliary items.

For a given formula $\varphi(x, y)$ with clauses C_1, \dots, C_L construct an instance $\text{CEx}_{\varphi(x, y)}$ of a combinatorial exchange with bidders and items as follows. First, consider $n + L + 2$ sellers:

- One seller S_i^X for each $i = 1, \dots, n$. Each seller S_i^X offers items χ_i and $\bar{\chi}_i$. These items will later indicate which logical values have to be assigned to the literals x such that $\forall y \varphi(x, y)$ is true.
- One seller S_l^ψ for each $l = 1, \dots, L$. Each seller S_l^ψ offers items $\psi_{lii'}$ for $i, i' = 1, \dots, n$. The sellers correspond to the clauses of $\varphi(x, y)$ and below we describe how an allocation of the items from a seller of type S^ψ corresponds to the truth value the corresponding clause evaluates to.
- One seller $S^{\gamma, \phi}$ who offers items $\gamma_j, \bar{\gamma}_j$ for $j = 1, \dots, m$ as well as items ϕ_{li} for $l = 1, \dots, L$ and $i = 1, \dots, n$. The items of type γ correspond to the possible values which can be assigned to literals y . The items ϕ_{li} are auxiliary items which indicate which clauses evaluate to false as a result of the assignment of y . While items of type ψ already correspond to the truth assignments of the clauses, these additional auxiliary items are necessary in the proof for stability reasons since seller $S^{\gamma, \phi}$ now also needs to be part of any blocking coalition involving items corresponding to the truth assignment of clauses.
- One seller S^λ who offers items λ_i^k and $\bar{\lambda}_i^k$ for $i = 1, \dots, n$ and $k = 1, 2$. These serve as auxiliary items to increase competition for buyers in order to drive up prices and deplete the budgets of buyers, as we will describe below


We introduce the following short notations for bundles of items:

- $\mathcal{T}_i^{\psi, \phi} = \{\psi_{lii'} | i' = 1, \dots, n\} \cup \{\phi_{li}\}$
- $\mathcal{F}_i^\psi = \{\psi_{lii'} | i' = 1, \dots, n\}$
- $\mathcal{T}_i^\psi = \{\psi_{lii'} | i, i' = 1, \dots, n\}$
- $\mathcal{F}_i^{\psi, \phi} = \mathcal{T}_i^\psi \cup \{\phi_{li} | i = 1, \dots, n\}$

Figure 2 illustrates an example for these bundles of items sold by S_1^ψ and $S^{\gamma, \phi}$. It can be seen that the bundles intersect with each other in such a way, that if for any $i \in \{1, \dots, n\}$, a bundle \mathcal{F}_i^ψ is purchased by a buyer, no bundle $\mathcal{T}_{i'}^{\psi, \phi}$ can be purchased for any $i' \in \{1, \dots, n\}$ and vice versa. Similarly, bundles $\mathcal{F}_i^{\psi, \phi}$ and \mathcal{T}_i^ψ intersect with all other bundles.

Next, we define the buyers with their preferences and budgets. Let $T < \frac{1}{n}$, $U > nL$, $V > 4U$ and $W > 7nV$. First, we define buyers of type B^K and type B^M whose assignments will directly correspond to the logical values of the literals x

- For $i = 1, \dots, n$ let B_i^K be a buyer with a budget of $V + T$ and a value of W for each of the following bundles:

Figure 2 Illustration of buyers' interests, only concerning items of type ψ 

- For $l = 1, \dots, L$, bundle $\mathcal{K}_l := \{\chi_i\} \cup \mathcal{T}_{li}^{\psi, \phi}$
- For $l = 1, \dots, L$, bundle $\bar{\mathcal{K}}_l := \{\bar{\chi}_i\} \cup \mathcal{T}_{li}^{\psi, \phi}$

He is interested in obtaining exactly one of these bundles and his value for obtaining one or more of the bundles is equal to the maximal value of his obtained bundles.

- For $i = 1, \dots, n$ let $B_i^{\mathcal{M}}$ be a buyer with a budget of $2V$ and a value of $2V$ for the bundles

- $\mathcal{M}_i := \{\chi_i, \lambda_i^1, \lambda_i^2\} \cup \bigcup_{x_i \in C_l} \mathcal{F}_{li}^{\psi}$
- $\bar{\mathcal{M}}_i := \{\bar{\chi}_i, \bar{\lambda}_i^1, \bar{\lambda}_i^2\} \cup \bigcup_{\bar{x}_i \in C_l} \mathcal{F}_{li}^{\psi}$

He is interested in exactly one of these bundles.

Buyers $B^{\mathcal{K}}$ and $B^{\mathcal{M}}$ are designed in such a way that for all $i \in \{1, \dots, n\}$, buyer $B_i^{\mathcal{K}}$ will obtain one of the items $\{\chi_i, \bar{\chi}_i\}$, while buyer $B_i^{\mathcal{M}}$ obtains the other item. Whenever $B_i^{\mathcal{K}}$ buys χ_i , this corresponds to an assignment of true to the corresponding x_i and whenever $B_i^{\mathcal{K}}$ buys $\bar{\chi}_i$ it corresponds to an assignment of false. Buyers $B_i^{\mathcal{M}}$ obtains the opposite item (corresponding to its negation) as well as the bundles \mathcal{F}_{li}^{ψ} for all l which evaluate to false due to the assignment of $B_i^{\mathcal{M}}$. The budgets and valuations are chosen in such a way, that none of the buyers described below can outbid buyers of type $B_i^{\mathcal{M}}$ at seller S_l^{ψ} , i.e. no bundles containing any item of \mathcal{T}_l^{ψ} can be sold when $B_i^{\mathcal{M}}$ desires \mathcal{F}_{li}^{ψ} for some $i \in \{1, \dots, n\}$. In the following, we say that $B_i^{\mathcal{M}}$ *blocks* the bundle \mathcal{T}_l^{ψ} (and therefore the bundles $\mathcal{T}_{li}^{\psi, \phi}$ for all i as well as bundle $\mathcal{F}_l^{\psi, \phi}$). We will see in the proof that in this case buyers of types $B^{\mathcal{G}}$ and $B^{\mathcal{K}}$ can only compete for unblocked bundles.

Additionally, we introduce the following auxiliary bidders who drive up prices in order to deplete the budgets of buyers $B^{\mathcal{K}}$ and $B^{\mathcal{M}}$.

- For $i = 1, \dots, n$, identical buyers $B_i^{\chi,1}$ and $B_i^{\chi,2}$ who are interested in one of χ_i or $\bar{\chi}_i$, have a valuation of V for both, as well as a budget of V
- For $i = 1, \dots, n$ one buyer $B_i^{\lambda,1}$ who has a budget of U and a value of V for bundle λ_i^1 and a value of $V - L$ for $\bar{\lambda}_i^1$
- For $i = 1, \dots, n$ one buyer $B_i^{\lambda,2}$ who has a budget of U and a value of $V - L$ for bundle λ_i^2 and a value of V for $\bar{\lambda}_i^2$

The reason for including these auxiliary buyers and items is to bind an amount of V of the budget of buyer $B_i^{\mathcal{K}}$ to purchase items from seller S^{λ} such that he only has a budget of T left to purchase his remaining

Bidder type	Value	Budget
B^K	W	$V + T$
B^M	$2V$	$2V$
B^G	1	L
$B^{\chi,1}, B^{\chi,2}$	V	V
$B^{\lambda,1}$	V for λ^1 and $V - L$ for $\bar{\lambda}^1$	U
$B^{\lambda,2}$	$V - L$ for λ^1 and V for $\bar{\lambda}^1$	U

Table 12 Values and budgets of key buyer types.

items from sellers of type S^ψ and from seller $S^{\gamma,\phi}$. In the following, we define the final set of buyers which compete with buyers B^M for these items.

- For $j = 1, \dots, m$ one buyer B_j^G who has a budget of L and a value of 1 for each bundle:

$$\begin{aligned} -\mathcal{G}_{jl} &= \{\gamma_j\} \cup \mathcal{F}_l^{\psi,\phi} \text{ for each } l = 1, \dots, L \text{ with } \bar{Y}_j \in C_l \\ -\bar{\mathcal{G}}_{lj} &= \{\bar{\gamma}_j\} \cup \mathcal{F}_l^{\psi,\phi} \text{ for each } l = 1, \dots, L \text{ with } Y_j \in C_l \end{aligned}$$

Define as $|\mathcal{G}_j|$ the number of bundles of type \mathcal{G} a buyer B_j^G obtains and with $|\bar{\mathcal{G}}_j|$ the number of bundles of type $\bar{\mathcal{G}}$ a buyer B_j^G obtains. Then, his valuation for obtaining a larger bundle containing one or more of the bundles defined above is equal to $\frac{\max\{|\mathcal{G}_j|, |\bar{\mathcal{G}}_j|\} - 1}{n(n+1)}$, i.e. the maximum number of bundles of type \mathcal{G} and type $\bar{\mathcal{G}}$ he obtains. Thus, each buyer B_j^G is only interested in obtaining bundles which do not include both items, γ_j and $\bar{\gamma}_j$. We refer to a bundle which includes only one of these items with a valuation of k as a *clean bundle of size k* .

These buyers are designed in a way such that they compete for all bundles $\mathcal{F}_l^{\psi,\phi}$ which are not blocked by a buyer of type B_i^M (buying bundle \mathcal{F}_{li}^ψ) since the latter has a larger budget and higher valuation and thus can always outbid buyers of type B^G . Whenever a buyer of type B^G obtains such a bundle, it corresponds to the corresponding clause to evaluate to false. Buyers of type B^G maximize their welfare by purchasing as many of these packages as possible, corresponding to causing as many clauses to evaluate to false as possible which are not do not already evaluate to false due to buyers of type B^M . Only if buyers of type B^G can not buy all of these bundles, buyers of type B^K can be assigned their bundles. Similarly to above, we say that a buyer of type B^G blocks bundle $\mathcal{T}_l^{\psi,\phi}$ for buyers of type B^K if he purchases a bundle which contains $\mathcal{F}_l^{\psi,\phi}$.

A.4. Reduction

In the following, we prove that there exists a core solution in $\text{CEx}_{\varphi(x,y)}$ with a social welfare of at least nW if and only if $\exists x \forall y \varphi(x,y)$ is true. We refer to such a core solution as an nW -equilibrium.

First, we will prove these auxiliary results:

Lemma 1: In an nW -equilibrium, for each $i = 1, \dots, n$, buyer B_i^K obtains one of the bundles he values at W .

Lemma 2: In an nW -equilibrium, for each $i = 1, \dots, n$, buyer B_i^K obtains one of the items χ_i^1 or $\bar{\chi}_i^1$, buyer B_i^M obtains the complementary item and both pay V to S_i^χ .

Lemma 3: In an nW -equilibrium, for $i = 1, \dots, n$ buyer B_i^M obtains all his required items from sellers S^ψ and S^λ .

Lemma 4: In any core allocation, buyers of type $B^{\mathcal{G}}$ maximize the combined size of their clean bundles among the ones not blocked by buyers of type $B^{\mathcal{M}}$.

Lemma 5: There is a nW -equilibrium if and only if for $j = 1, \dots, m$, buyers $B_j^{\mathcal{G}}$ are not able to block all the remaining bundles for buyers of type $B^{\mathcal{K}}$.

Using these auxiliary results, we will be able to prove the main result.

LEMMA 1. *In an nW -equilibrium, for each $i = 1, \dots, n$, buyer $B_i^{\mathcal{K}}$ obtains one of the bundles he values at W .*

Proof: Assume that a buyer $B_i^{\mathcal{K}}$ does not obtain his preferred bundle (and thus, the total welfare generated by the other buyers $B_i^{\mathcal{K}}$ for $i' \neq i$ is at most $(n-1)W$). Then, there is no way to achieve a social welfare of at least nW since

$$W > 7nV > \underbrace{2nV}_{\text{Buyers } B^{\mathcal{M}}} + \underbrace{2nV}_{\text{Buyers } B^{\mathcal{X}}} + \underbrace{2nV}_{\text{Buyers } B^{\mathcal{A}}} + \underbrace{nL}_{\text{Buyers } B^{\mathcal{G}}}$$

which is an upper bound on the welfare achievable by all other buyers. Q.E.D.

Thus, in an nW -equilibrium, all buyers of type $B^{\mathcal{K}}$ obtain one of their desired bundles. As we described in the transformation, this is only possible, if there exists at least one $l \in \{1, \dots, L\}$, for which neither buyers of type $B^{\mathcal{M}}$ nor of type $B^{\mathcal{G}}$ block the bundle \mathcal{T}_l^{ψ} .

The following lemma is a simple observation how auxiliary buyers of type $B^{\mathcal{X}}$ are used to deplete the budget of buyers of type $B^{\mathcal{K}}$.

LEMMA 2. *In an nW -equilibrium, for each $i = 1, \dots, n$, buyer $B_i^{\mathcal{K}}$ obtains one of the items χ_i^1 or $\bar{\chi}_i^1$, buyer $B_i^{\mathcal{M}}$ obtains the other item and both pay V to $S_i^{\mathcal{X}}$.*

Proof: If either $B_i^{\mathcal{K}}$ or $B_i^{\mathcal{M}}$ would pay less than V , then either $B_i^{\mathcal{X},1}$ or $B_i^{\mathcal{X},2}$ could outbid them and obtain the respective items: In this case, seller $S_i^{\mathcal{X}}$, all buyers which obtain items from $S_i^{\mathcal{X}}$ (and in consequence all further buyers and sellers) can form a coalition and share the additional payment of the buyer of type $B^{\mathcal{X}}$ such that all members of this coalitions improve their payoffs. Thus, an assignment where $B_i^{\mathcal{K}}$ obtains an item from $S_i^{\mathcal{X}}$ but pays less than V can not be in the core. Since all buyers of type $B^{\mathcal{K}}$ need to obtain one of these items in order to reach an nW -equilibrium, the Lemma holds. Q.E.D.

LEMMA 3. *In an nW -equilibrium, for $i = 1, \dots, n$ buyer $B_i^{\mathcal{M}}$ obtains all his required items from sellers S^{ψ} and S^{λ} .*

Proof: Because of Lemma 2, in an nW -equilibrium, $B_i^{\mathcal{M}}$ needs to pay V for the item he obtains from seller $S_i^{\mathcal{X}}$. Then, he has a budget of V left to obtain the missing items from seller S^{λ} and sellers S_i^{ψ} in order to complete his desired bundle. He needs to purchase items of the form \mathcal{F}_{li}^{ψ} from S_i^{ψ} as well as either $\{\lambda_i^1, \lambda_i^2\}$ or $\{\bar{\lambda}_i^1, \bar{\lambda}_i^2\}$. No other buyer $B_j^{\mathcal{M}}$ with $j \neq i$ is interested in obtaining any of these items since they appear in no bundles with positive valuation for them. The only buyers interested in a subset these items are buyers $B_i^{\mathcal{K}}$ for $i' = 1, \dots, n$ (who only have a budget of T left due to Lemma 2), buyers of type $B^{\mathcal{G}}$ (who have a budget of at most L each) and buyers $B_i^{\lambda,1}$ and $B_i^{\lambda,2}$ (with a budget of U each). Since

$$V > 4U > \underbrace{T}_{\text{Buyers } B^{\mathcal{M}}} + \underbrace{nL}_{\text{Buyers } B^{\mathcal{G}}} + \underbrace{2U}_{\text{Buyers } B^{\lambda}},$$

buyer B_i^M can pay sellers S^ψ and S^λ enough to obtain his required items and there is no combination of buyers that can outbid B_i^M in order to form a coalition with the sellers such that all improve. Q.E.D.

The previous Lemma 3 showed that for any $i = 1, \dots, n$, buyer B_i^M gets all the items he requires from sellers S^ψ and S^λ and in particular all his required bundles of the form \mathcal{F}_i^ψ . Thus, he blocks the bundle \mathcal{T}_i and therefore also all bundles $\mathcal{F}_i^{\psi, \phi}$ for sellers of type B^G .

LEMMA 4. *In any core allocation, buyers of type B^G maximize the combined size of their clean bundles among the ones not blocked by buyers of type B^M .*

Proof: Assume that the maximum combined size of non-blocked clean bundles which can be obtained by buyers B^G is K but that in the core solution buyers only buy clean bundles with a combined size of $\kappa \leq K - 1$. There are no other buyers except for those of type B^K which are interested in any of the items offered by sellers S^λ or sellers S_l^ψ for those l for which \mathcal{T}_l is not blocked. Since $1 > nT$, there can be a coalition of those sellers and buyers B^G which can generate a value of $K > \kappa$ and distribute the welfare such that all participants are better off. This is a contradiction to the allocation being in the core. Then, if all buyers B^G pay the valuation of their obtained bundle to seller S^λ , there is no coalition among these sellers and buyers which want to deviate since S^λ can never improve upon his payoff. Q.E.D.

LEMMA 5. *There is a nW -equilibrium if and only if for $j = 1, \dots, m$, buyers B_j^G are not able to block all the remaining bundles for buyers of type B^K .*

Proof: For any i, l , buyer B_i^K is only able to obtain one of the sets $\mathcal{T}_i^{\psi, \phi}$ if it is neither blocked by a buyer B^M or B^G . Thus, there is some l for which all buyers can obtain these items if and only if buyers B^G do not block all of these bundles and as of Lemma 1 there is an nW -equilibrium if and only if all buyers B^K obtain one of their bundles' values at W Q.E.D.

THEOREM 2. *There exists an nW -equilibrium if and only if $\exists x \forall y \varphi(x, y)$ is true.*

Proof: Consider an nW -equilibrium and set x_i to true if buyer B_i^K obtains item χ_i and set x_i to false if he obtains $\bar{\chi}_i$. Then, buyer B_i^M obtains the negated item and bundles \mathcal{F}_i^ψ for all clauses C_l which evaluate to false due to the assignment of x_i . This is equivalent to blocking the bundles \mathcal{T}_l for buyers B^G who thus compete for the non-blocked bundles. Each combination of bundles obtained by buyer B_j^G resembles a number of clauses which can be made false by a truth assignment of y_j . If B_j^G obtains γ_j this corresponds to an assignment of y_j to true and if he obtains $\bar{\gamma}_j$ it corresponds to an assignment of y_j to false. As of Lemma 3 and 4, in any core allocation (and hence, especially in an nW -equilibrium) buyers B^G try to maximize their combined number of bundles not blocked by buyers of type B^M which is equivalent to blocking as many bundles as possible for buyers of type B^K . This corresponds to assigning truth values to y so that as many clauses as possible evaluate to false in φ . However, since by assumption the assignment results in an nW -equilibrium, buyers B^G are not successful in blocking all bundles because of Lemma 5. Therefore, there is no assignment of variables y such that $\varphi(x, y)$ can be set to false for this assignment of x .

Conversely, let x be a truth assignment such that $\forall y \varphi(x, y)$ is true. Then, consider the following trades in the combinatorial exchange: Trades for seller S_i^X for $i = 1, \dots, n$:

- For $i = 1, \dots, n$, if x_i is true, assign to buyer B_i^K items χ_i for a price of V .
- For $i = 1, \dots, n$, if x_i is false, assign to buyer B_i^K items $\bar{\chi}_i$ for a price of V .
- For $i = 1, \dots, n$, assign to buyer B_i^M the item not allocated to B_i^K for a price of V .

Trades for seller S^λ :

- For $i = 1, \dots, n$, if x_i is true, assign to buyer B_i^M the items $\bar{\lambda}_i^1, \bar{\lambda}_i^2$ for a price of $2U$, as well as item λ_i^1 to $B_i^{\lambda,1}$ and λ_i^2 to $B_i^{\lambda,2}$ for a price of U each.
- For $i = 1, \dots, n$, if x_i is false, assign to buyer B_i^M the items λ_i^1, λ_i^2 for a price of $2U$, as well as item $\bar{\lambda}_i^1$ to $B_i^{\lambda,1}$ and $\bar{\lambda}_i^2$ to $B_i^{\lambda,2}$ for a price of U each.

Further, assign to buyers B_i^M his remaining required items from sellers S^ψ , paying a price of 1 to each seller he purchases from. Then, there is a nW -equilibrium which extends these assignments. Similar to the first part of the proof, a buyer which blocks a bundle \mathcal{T}_l for the other buyers corresponds to a truth assignment of the corresponding variable which results in the clause l to evaluate to false. Since $\bar{A}y : \neg\varphi(x, y)$ for the truth assignment of x , buyers B^M and B^G can not block all bundles for buyers B^K , so each of them can obtain a bundle which he values at W . There is no coalition of buyers and sellers which want to deviate from this equilibrium:

- Buyers B^G , sellers S^ψ and $S^{\gamma,\phi}$ can't form a coalition exclusively among themselves because of Lemma 5.
- For all $i = 1, \dots, n$, there exists no coalition including buyers B_i^M in which all participants can be made better off: Since B_i^M needs to pay sellers S_i^X and S^λ more money in order for them to join the coalition, he needs to pay less to the sellers S^ψ he switches to. Those sellers are disjoint from the sellers S^ψ he purchased from earlier. Thus, he can save at most $\frac{L}{2}$ units from switching which he needs to redistribute to S_i^X and S^λ . However, since buyers $B_i^{\lambda,1}$ and $B_i^{\lambda,2}$ are affected by these trades as well, they need to be in the coalition as well and purchase items such that S^λ can be made better off (since $\frac{L}{2} < 2D$, buyer S^λ can not deviate only with B_i^M). However, since for one of the two buyers, his new payoff is reduced by L , he will not agree to this coalition unless his payment is also reduced by at least L . However, since the second of these two buyers can not pay more as he is already capped by his budget, this is not possible.
- All other buyers and sellers can not deviate from the grand coalition on their own but need at least one buyer B^M in order for all members to achieve a higher payoff. As by the above, there is no coalition including a buyer B^M that can achieve this.

Thus, for a given truth assignment, there is a nW -equilibrium and the proof is complete. Q.E.D.