Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)

Core-Selecting Payment Rules for Combinatorial Auctions with Uncertain Availability of Goods

Dmitry Moor and **Sven Seuken** and **Tobias Grubenmann** and **Abraham Bernstein**
Department of Informatics
University of Zurich
Zurich, Switzerland
{dmitry.moor, seuken, grubenmann, bernstein}@ifi.uzh.ch

Abstract

In some auction domains, there is uncertainty regarding the final availability of the goods being auctioned off. For example, a government may auction off spectrum from its public safety network, but it may need this spectrum back in times of emergency. In such a domain, standard combinatorial auctions perform poorly because they lead to violations of individual rationality (IR), even in expectation, and to very low efficiency. In this paper, we study the design of core-selecting payment rules for such domains. Surprisingly, we show that in this new domain, there does not exist a payment rule with is guaranteed to be ex-post core-selecting. However, we show that by designing rules that are "execution-contingent," i.e., by charging payments that are conditioned on the realization of the availability of the goods, we can reduce IR violations. We design two core-selecting rules that always satisfy IR in expectation. To study the performance of our rules we perform a computational Bayes-Nash equilibrium analysis. We show that, in equilibrium, our new rules have better incentives, higher efficiency, and a lower rate of ex-post IR violations than standard core-selecting rules.

1 Introduction

Combinatorial auctions (CAs) have been successfully applied in many real-world settings, including procurement auctions [Sandholm, 2013], TV advertising auctions [Goetzendorf *et al.*, 2015], and government spectrum auctions [Cramton, 2013; Ausubel and Baranov, 2014]. CAs are specifically designed for domains where bidders can have complex preferences over bundles of heterogeneous items. An important and seemingly "innocent" assumption in auction design is that all of the goods will be available for consumption by the winning bidders. While this assumption is often satisfied, there are some domains where this is not the case. In these domains, standard mechanisms perform poorly and new designs are required.

1.1 Uncertain Availability: A Motivating Example

In the US, public safety networks are used by the police, firefighters, and emergency medical technicians during times of emergencies. In 2012, following the events of 9/11 and hurricane Katrina, the US congress even reserved some parts of the 700MHz spectrum "to public safety for use in a nationwide broadband network" [FCC, 2016]. However, this legislation also allows for the use of this spectrum by private companies when the spectrum would otherwise be idle. The company *Rivada Networks* for example is currently designing an auction platform for this purpose [Cramton and Doyle, 2015]. These auctions can happen on a weekly or daily basis, or even in real-time; but bidders have to accept the risk that the spectrum they purchased in the auction might become unavailable because an emergency occured and the spectrum is needed for public safety reasons.

There are many other CA domains with this kind of *uncertain availability of goods*. For example, the company *BandwidthX* has developed a platform to auction off bandwidth from wireless hotspots, but those hotspots can also become unavailable at any point in time [Seuken *et al.*, 2015]. Furthermore, Moor et al. [2015] introduced an auction for data, where the sellers only have imprecise estimates of the data they will actually be able to supply.

1.2 Execution-Contingent Auctions

Domains with uncertain availability of goods pose new challenges. As we will show, standard mechanisms perform very poorly: they violate individual rationality (IR), even in expectation, they have bad incentives, and low efficiency in equilibrium. To address these shortcomings, Porter et al. [2008] introduced *execution-contingent* mechanisms, which charge payments dependent on the realized availabilities of the goods (see [Ceppi *et al.*, 2015; Ramchurn *et al.*, 2009] for extensions). Using this paradigm, one can design mechanisms that satisfy IR and strategyproofness *in expectation*. For domains as described above, where many small auctions may be run repeatedly over time, *some* IR violations may be acceptable, as long as IR in expectation is satisfied.

However, the execution-contingent mechanisms studied in the literature so far [Porter *et al.*, 2008; Ceppi *et al.*, 2015; Ramchurn *et al.*, 2009] are not suitable for CAs in practice because they charge VCG-like payments. Unfortunately, in a CA domain, VCG has numerous problems [Ausubel and Milgrom, 2006]. Most notably, VCG can lead to very low or even zero revenue, which opens up opportunities for collusion between the seller and collections of bidders.

1.3 Core-Selecting Payment Rules

These drawbacks of VCG have led to the development of *core-selecting payment rules* which offer a principled way to ensure that the revenue in the auction is high enough such that there are no opportunities for collusion [Milgrom, 2007; Day and Milgrom, 2008; Day and Raghavan, 2007; Ausubel and Baranov, 2016]. They have already been used successfully in practice, for example in the UK, Canada, the Netherlands, and Switzerland to auction off billions of dollars worth of 4G spectrum. In this paper, we design *execution-contingent core-selecting payment rules* for domains with uncertainty availability of goods. The first question we ask is "what is the right notion of an execution-contingent core?" Second, we find that our execution-contingent cores may sometimes be empty, and thus, our rules will only be "core-selecting" in a relaxed sense. Finally, in contrast to prior work on execution-contingent mechanisms, we are not fully content with rules that satisfy "IR in expectation," but also want to minimize the rate of ex-post IR violations.

1.4 Overview of Contributions

Our goal in this paper is to design payment rules that work well in a domain with uncertain availability of goods. We make the following contributions:

1. We generalize the EC-VCG mechanism introduced by Porter et al. [2008] to domains with *continuous* and *dependent* availabilities of goods.
2. We introduce two execution-contingent core-selecting payment rules which satisfy IR in expectation.
3. We perform a computational Bayes-Nash equilibrium (BNE) analysis evaluating our new rules in terms of incentives, efficiency, and ex-post IR violations.

The results from our BNE analysis show that our new payment rules have better incentives, higher efficiency, and a lower rate of ex-post IR violations than standard core-selecting rules.

2 Preliminaries

2.1 Formal Model

Let $N = \{1, ..., n\}$ be a set of bidders and let s be the seller. Let $G = \{A, B, C, ...\}$ be the set of goods, with $|G| = m$. Each bidder has a valuation function $v_i : 2^G \to \mathbb{R}$ which specifies i's value for every possible bundle $S \in 2^G$.

We consider a *two time-period* model: at *allocation time*, there is uncertainty regarding which bundles will eventually be available, and at *consumption time*, some of the availability (depending on the mechanism) will be revealed. With every bundle S we associate a random variable $a(S)$ which represents the *availability* of the bundle S, i.e., the extent to which S will be available at the time of consumption. We let $\tilde{a}(S)$ denote the realization of the random variable $a(S)$ at consumption time.

Let f be a joint probability mass function for all random variables $a(S)$. We assume that f is exogenous and known by all agents.

The seller may have costs for providing the goods. We denote costs by $\{c_1, ..., c_m\}$, where c_j is the cost of the j-th good. We assume that the cost function $c(S) \in \mathbb{R}$ is additive, such that $c(S) = \sum_S c_j$. We assume that the seller does not strategize, as is common in the analysis of CAs [Day and Cramton, 2012].

An allocation is denoted as $x = \{x_1, ..., x_n\}$, with x_i denoting the bundle allocated to agent i and x_{-i} denoting the vector of bundles of all agents except i. Similarly, for any $K \subset N$, we use x_K when referring to an allocation among all bidders in K and use x_{-K} to denote an allocation among bidders in $N \backslash K$. We let $v_i(x_i)$ denote bidder i's true value for its allocated bundle x_i, and we let $\hat{v}_i(x_i)$ denote bidder i's value report for bundle x_i (possibly non-truthful). We let $a(x_i)$ denote the random variable corresponding to bundle x_i, and $\tilde{a}(x_i^*)$ denote the *realized availability* corresponding to x_i^*. If bidder i is allocated bundle x_i^* then his *realized value* is thus given by $v_i(x_i^*)\tilde{a}(x_i^*)$.

Assumption 1. *We assume that all bidders are* "extremely" single-minded*, i.e., each bidder has non-zero value for exactly one bundle $S \subseteq 2^G$, and zero value for all other bundles $S' \neq S$ (including supersets of S).*

Remark 1. *This assumption allows for the simple definition of the realized value we have just provided, i.e., $v_i(x_i^*)\tilde{a}(x_i^*)$. Without this assumption, we would have to consider all sub-bundles of a bidder's allocated bundle to compute his realized value, i.e., $\max_{S \subseteq x_i^*}\{v(S)\tilde{a}(S)\}$, which may be computationally infeasible in domains with a large number of goods. But more importantly, we make this assumption to simplify the notation and the analysis of the mechanisms we will present. In future work, we will extend our results to the full domain.*

Let x be an allocation. We let W_x (the *winners*) denote the set of allocated agents under this allocation, i.e., $W_x = \{i | x_i \neq \emptyset\}$. The social welfare of the allocation x is

$$SW(x) = \sum_{i \in W_x} (v_i(x_i) - c(x_i))\tilde{a}(x_i).$$

We assume quasilinear utilities $u_i(x_i, p_i) = v_i(x_i) - p_i$, where p_i is bidder i's payment for x_i; we let p_s denote the payment received by the seller. We let $p = (p_s, p_1, p_2, ..., p_n)$ denote the vector of payments received by the seller and paid by all bidders. We let $O = \langle x, p \rangle$ denote an outcome, i.e., an allocation and the payment vector.

2.2 Properties of Mechanisms

Let $v = (v_1, ..., v_n)$ and $c = (c_1, ..., c_m)$. We let $\tilde{a}$ denote the vector of availabilities at consumption time (which may depend on the mechanism and the domain), i.e., $\tilde{a} = (\tilde{a}(S) : \forall S \in 2^G, \tilde{a}(S)$ is known at consumption time). We let $\mathcal{M} = \langle g, h \rangle$ denote a mechanism, where $g(v, c, f) = x$ is an allocation rule and $h(x, v, c, f, \tilde{a}_M) = p$ is a payment rule. We now define a number of standard mechanism design properties; however, because we consider a domain with uncertainty, we need to define most of these properties "in expectation."

Definition 1. *A mechanism $\mathcal{M} = \langle g, h \rangle$ is* ***strategyproof in expectation*** *if $\forall i \in N$, $\forall v_i$, for all $\hat{v}_{-i}$*

$$\mathbb{E}_f[u_i(g(v_i, \hat{v}_{-i}, c, f), p_i)] \geq \mathbb{E}_f[u_i(g(\hat{v}_i, \hat{v}_{-i}, c, f), p_i)].$$

Definition 2. *A mechanism* $\mathcal{M} = \langle g, h \rangle$ *with* $h(g(v,c,f), v, c, f, \tilde{a}_M) = p$ *is ex-post* ***individually rational (IR)*** *if* $\forall i \in N$, $\forall v_i, \forall \hat{v}_{-i}$

$$v_i(g(v_i, \hat{v}_{-i}, c, f))\tilde{a}(g(v_i, \hat{v}_{-i}, c, f)) - p_i \geq 0.$$

Definition 3. *A mechanism* $\mathcal{M} = \langle g, h \rangle$ *with* $h(g(v,c,f), v, c, f, \tilde{a}_M) = p$ *is* ***individually rational in expectation (IRE)*** *if* $\forall i \in N$, $\forall v_i$, *for all* $\hat{v}_{-i}$

$$\mathbb{E}_f[u_i(g(v_i, \hat{v}_{-i}, c, f), p_i)] \geq 0.$$

Definition 4. *The* ***rate of ex-post IR violations*** *of a mechanism* $\mathcal{M} = \langle g, h \rangle$ *with* $h(g(v,c,f), v, c, f, \tilde{a}_M) = p$ *is defined as the following probability:*

$$\mathbb{P}(v_i(g(v_i, \hat{v}_{-i}, c, f))\tilde{a}(g(v_i, \hat{v}_{-i}, c, f)) - p_i < 0) \quad (1)$$

Definition 5. *A mechanism is* ***budget balanced*** *if the sum of all payments paid by the bidders is equal to the payment received by the seller, i.e.,*

$$\sum_{i \in N} p_i = p_s.$$

Definition 6. *A mechanism is* ***expected social welfare maximizing*** *if its allocation rule selects an allocation* x *with*

$$x \in \operatorname*{argmax}_x \mathbb{E}_f[SW(x)].$$

2.3 VCG Mechanism

The famous VCG mechanism [Vickrey, 1961; Clarke, 1971; Groves, 1973] selects a social welfare maximizing allocation and computes payments equal to the externality each agent imposes on all other agents. We let x^* denote the allocation which maximizes social welfare when all agents are considered, and x^{-i} denotes the allocation which maximizes social social welfare when all agents except i are considered. VCG payments are then:

$$p_i^{\text{VCG}} = SW(x^{-i}) - SW_{-i}(x^*).$$

VCG is a particularly attractive mechanism because, in a domain without uncertainty about the availability of goods, it is social welfare maximizing, strategyproof and satisfies IR.

3 Execution-Contingent VCG

Porter et al. [2008] generalized the VCG mechanism to domains with uncertain availability of goods by introducing an *execution-contingent* variant of VCG. The main idea is to make payments contingent on the realized availabilities, which implies that the payments are not computed at allocation time, but at consumption time (see Figure 1). They considered a domain with binary and independent random variables capturing the availabilities, and proved that in this domain, their mechanism is strategyproof and IR in expectation. However, in the domains we described in the beginning, the availabilities of the goods will typically be *dependent* (e.g., consider a terrorist attack affecting a whole city) and the availabilities may be continuous (e.g., a resource can be used partially). We now introduce the ECC-VCG mechanism, which generalizes the mechanism introduced by Porter et al. [2008] to also handle continuous, dependent random variables.

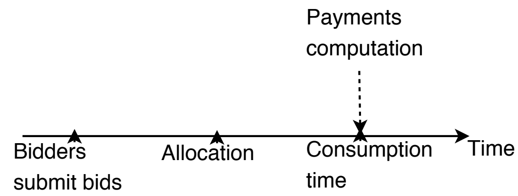

Figure 1: Flow chart of execution-contingent mechanisms.

ECC-VCG Mechanism.

- *Allocation rule*: select $x^* \in \operatorname{argmax}_x \mathbb{E}_f[SW(x)]$
- *Payment rule*:

$$p_i^{\text{ECC-VCG}} = \mathbb{E}_f[SW(x^{-i})|\tilde{a}(x_j^*), j \in W_{x^*}] - SW_{-i}(x^*)$$

The idea behind this *execution-contingent conditional VCG* mechanism is similar to standard VCG: $p_i^{\text{ECC-VCG}}$ represents a bidder's externality on all other bidders, except that we do not know this externality exactly and thus compute an estimate by taking the conditional expectation. The following example illustrates the ECC-VCG mechanism.

Example 1. *Consider a setting with three bidders* $N = \{1, 2, 3\}$*, two goods* $G = \{A, B\}$ *and a seller* s *with no costs. The bidders' values are provided in the following table:*

	A	B	$\{A, B\}$
Bidder 1	*0.1*		
Bidder 2		*0.2*	
Bidder 3			*0.3*

Let $a(\{A\}), a(\{B\}), a(\{A, B\})$ *denote availabilities of* $\{A\}, \{B\}$ *and* $\{A, B\}$ *respectively; let* f *denote the joint probability mass function for these random variables. The following table specifies* f*:*

(a) $a_1 = 0$

a_2	$a_3 = 0$	$a_3 = 1$
0	0.25	0
1	0.25	0

(b) $a_1 = 1$

a_2	$a_3 = 0$	$a_3 = 1$
0	0.1	0
1	0	0.4

It is easy verify that the expected marginal availabilities of the bundles $\{A\}$, $\{B\}$*, and* $\{A, B\}$ *are 0.5, 0.65 and 0.4 respectively. Thus, the efficient allocation is to allocate* A *to bidder* 1 *and* B *to bidder* 2*, and* $\mathbb{E}_f[SW(x)] = 0.1 \cdot 0.5 + 0.2 \cdot 0.65 = 0.18$.

Now assume that after the allocation has happened the realized availabilities are $\tilde{a}_1 = 0$ *and* $\tilde{a}_2 = 1$*. Thus* $\mathbb{E}_f[a_3|\tilde{a}_1, \tilde{a}_2] = 0$*. The execution-contingent payments are:*

$$p_1^{ECC\text{-}VCG} = 0.2 \cdot \tilde{a}_2 - (0.2 \cdot \tilde{a}_2) = 0 \quad (2)$$

$$p_2^{ECC\text{-}VCG} = 0 - (0.1 \cdot \tilde{a}_1) = 0 \quad (3)$$

However, if realized availabilities were $\tilde{a}_1 = 1$ *and* $\tilde{a}_2 = 1$*, then* $\mathbb{E}_f[a_3|\tilde{a}_1, \tilde{a}_2] = 0.4$ *and execution contingent payments would be*

$$p_1^{ECC\text{-}VCG} = 0.2 \cdot \tilde{a}_2 - (0.2 \cdot \tilde{a}_2) = 0 \quad (4)$$

$$p_2^{ECC\text{-}VCG} = 0.3 \cdot \mathbb{E}_f[a_3|\tilde{a}_1, \tilde{a}_2] - (0.1 \cdot \tilde{a}_1) = 0.02. \quad (5)$$

Note that the ECC-VCG mechanism only takes the realization of the *allocated* bundles into account when computing payments. If we know the realizations of *all* bundles at consumption time, then we can use the following mechanism:

ECR-VCG Mechanism.

- *Allocation rule*: select $x^* \in \text{argmax}_x \mathbb{E}_f[SW(x)]$
- *Payment rule*:

$$p_i^{\text{ECR-VCG}} = \mathbb{E}_f[SW(x^{-i})|\tilde{a}(S), \forall S] - SW_{-i}(x^*)$$

This *execution-contingent realized VCG* mechanism is a special case of ECC-VCG, as the payment rule uses a more precise estimate of the externality imposed on other bidders.

Theorem 1. *ECC-VCG and ECR-VCG are strategyproof in expectation.*

Proof. We prove the theorem for ECC-VCG; the proof for ECR-VCG is analogous. We let $\hat{v}_i$ denote a reported value function of agent $i \in N$ and v_i is the true value of the agent. Let x^* be the optimal allocation given $(\hat{v}_i, \hat{v}_{-i})$. Consider the conditional expectation in the payment rule:

$$\mathbb{E}_f[SW(x^{-i})|\tilde{a}(x_j^*), j \in W_{x^*}] \quad (6)$$

$$= \sum_{k \in W_{x^{-i}}} (\hat{v}_k(x_k^{-i}) - c(x_k^{-i}))\mathbb{E}_f[a(x_k^{-i})|\tilde{a}(x_j^*), j \in W_{x^*}]. \quad (7)$$

Now the payment of a bidder $i \in W_{x^*}$ is

$$p_i^{\text{ECC-VCG}} = \sum_{k \in W_{x^{-i}}} (\hat{v}_k(x_k^{-i}) - c(x_k^{-i}))\mathbb{E}_f[a(x_k^{-i})|\tilde{a}(x_j^*), j \in W_{x^*}]$$
$$- \sum_{k \in W_{x^*} \setminus i} (\hat{v}_k(x_k^*) - c(x_k^*))\tilde{a}(x_k^*) + c(x_i^*)\tilde{a}(x_i^*)$$

The expected utility of bidder i in this case is

$$\mathbb{E}_f[u_i(x^*)] = \mathbb{E}_f[v_i(x_i^*)a(x_i^*) - p_i^{\text{ECC-VCG}}] =$$
$$\mathbb{E}_f\Big[v_i(x_i^*)a(x_i^*) - c(x_i^*)a(x_i^*) + \sum_{k \in W_{x^*} \setminus i} (\hat{v}_k(x_k^*) - c(x_k^*))a(x_k^*)$$
$$\underbrace{- \sum_{k \in W_{x^{-i}}} (\hat{v}_k(x_k^{-i}) - c(x_k^{-i}))\mathbb{E}_f[a(x_k^{-i})|a(x_j^*), j \in W_{x^*}]}_{\texttt{does not depend on } \hat{v}_i}\Big]$$

The first three terms under the last expectation is equal to expected social welfare if the agent reports truthfully. Thus, because the allocation rule maximizes expected social welfare, and because the last expression is independent of i's report, truthful reporting maximizes agent i's expected utility. In this case: $\mathbb{E}_f[u_i(x^*)] = \mathbb{E}_f[SW(x^*)] - \mathbb{E}_f[SW(x^{-i})]$. □

Theorem 2. *ECC-VCG and ECR-VCG are individually rational in expectation.*

Proof. We show the theorem for ECC-VCG. From the proof of Theorem 1 we know that under truthful reporting

$$\mathbb{E}_f[u_i(x^*)] = \mathbb{E}_f[SW(x^*)] - \mathbb{E}_f[SW(x^{-i})]. \quad (8)$$

Because the allocation rule selects an expected social welfare maximizing allocation we know that $\mathbb{E}_f[SW(x^*)] \geq \mathbb{E}_f[SW(x^{-i})]$), and thus $\mathbb{E}_f[u_i(x^*)] \geq 0$. The proof is analogous for ECR-VCG. □

4 Core-Selecting Payment Rules

VCG-based rules are attractive because they are strategyproof (or strategyproof in expectation, in our domain). However, they are not suitable in CAs because they can lead to very low or even zero revenue, which opens up opportunities for collusion between the seller and collections of bidders. In this section, we introduce our *execution-contingent core-selecting payment rules*. To this end, we first need some definitions:

Definition 7. *An outcome O is blocked by a coalition $K \subset N$ of bidders, if there exists another outcome $\bar{O}$ which is weakly preferred over O by every bidder $i \in K$ and which provides higher utility for the seller. In this case K is called a blocking coalition.*

Definition 8. *An outcome O is in the core if it is (a) individually rational and (b) is not blocked by any coalition.*

If $O^* = \langle x^*, p\rangle$ and $K \subset N$ is a coalition of bidders, then $\sum_{i \in K}(v_i(x_i^*) - p_i)$ is the total opportunity cost of agents in the coalition. We let $\omega(K)$ denote the total welfare which the coalition could achieve by redistributing goods among themselves. The maximum additional value which K could give to the seller is $\omega(K) - \sum_{i \in K}(v_i(x_i^*) - p_i)$. A core constraint guarantees that this amount is not larger than what the seller can already get under the current allocation, and this core constraint has to hold for all possible coalitions:

$$\sum_{i \in W_{x^*}} (p_i - c(x_i^*)) \geq \omega(K) - \sum_{i \in K}(v_i(x_i^*) - p_i) \quad \forall K \subseteq N$$

A core-selecting mechanism selects a social welfare maximizing allocation and picks payments from the core. Unfortunately, in general combinatorial auction domains (with complements), even without uncertain availability of goods, there does not exist a payment rule that is Bayes-Nash incentive compatible and core-selecting [Goeree and Lien, 2016]. Of course, this impossibility extends to our new domain. Thus, none of our core-selecting payment rules can be truthful.

The core in a domain with uncertainty. In contrast to the standard domain, defining the core in a domain with uncertain availability of goods is less straightforward. The difficulty arises from the fact that in a domain with uncertainty agents might be willing to deviate either before or after availabilities of goods are realized. In the former case, we can talk about an *ex-ante* core while the latter case implies an *ex-post* core. Although these two concepts differ only in the amount of information which is used to evaluate the expected total welfare achieved by a coalition, we will show that these concepts actually lead to quite different properties.

In the next section, we will first present an ex-ante core-selecting payment rule, which is not execution contingent. However, as we will later see in Section 5, this leads to a relatively large rate of IR violations, compared to execution-contingent payment rules. To address this, we will then study ex-post core-selecting payment rules in Sections 4.2-4.5.

4.1 Ex-ante Core-Selecting Payment rules

A straightforward approach to generalize the idea of a core to a domain with uncertainty is to apply the core constraints as defined in the previous section to the expected values of the bidders. In this case the total expected welfare $\mathbb{E}_f[\omega(K)]$ that a coalition K can achieve in expectation is $\mathbb{E}_f[\omega(K)] = \max_x \mathbb{E}_f[\sum_{i\in K}(v_i(x) - c_i(x))a(x)]$. By limiting this value to be smaller than the expected total utility which agents in the coalition can get under the current allocation we ensure that the agents in the coalition are not be willing to deviate from the current allocation. Formally, this gives rise to the following set of core constraints, which have to hold for all $K \subseteq N$:

$$\mathbb{E}_f\Big[\sum_{i\in W_{x^*}} (p_i - c(x_i^*)a(x_i^*))\Big] \geq \tag{9}$$

$$\mathbb{E}_f\Big[\omega(K)\Big] - \sum_{i\in K}\mathbb{E}_f\Big[(v_i(x_i^*)a(x_i^*) - p_i)\Big] \tag{10}$$

Even though this core is not execution contingent, it provides all core properties ex-ante, i.e., before availabilities of goods are realized. Furthermore, the following corollary says that this core is never empty:

Corollary 1. *In a domain with uncertain availability of goods, the ex-ante core is never empty.*

The corollary follows from the fact that the core must contain at least one point, namely pay-as-bid (see [Day and Milgrom, 2008, Footnote 1]).

4.2 Impossibility of Ex-post Core-Selecting Payment Rules in Domains with Uncertainty

In this section, we will show the surprising result that there does not, in general, exist an ex-post core-selecting payment rule in a domain with uncertain availability of goods. For this, we first need a few more definitions. Assume that $\langle x^*, p\rangle$ is an auction outcome and let $L \subseteq 2^G$ be a fixed set of bundles.

Definition 9. *A generalized expected coalitional value $\omega_{ge}(K, L)$ of a coalition $K \subseteq N$ given the set L is the maximum expected welfare the coalition can achieve given realized availabilities corresponding to bundles from L. Formally,*

$$\omega_{ge}(K,L) = \max_x \mathbb{E}_f\Big[\sum_{i\in K}(v_i(x_i) - c(x_i))a(x_i)\Big|\tilde{a}(S), S\in L\Big]$$

Note that bidders have a total opportunity cost of $\sum_{i\in K}(v_i(x_i^*)\tilde{a}(x_i^*) - p_i)$ for joining coalition K. If they decide to join, then in expectation they can achieve a total welfare of $\omega_{ge}(K, L)$ and thus they can provide at most $\omega_{ge}(K, L) - \sum_{i\in K}(v_i(x_i^*)\tilde{a}(x_i^*) - p_i)$ of additional value to the seller, which gives rise to the following set of core constraints:

Generalized Ex-post Core Constraint $\forall K \subseteq N$:

$$\sum_{i\in W_{x^*}} (p_i - c(x_i^*)\tilde{a}(x_i^*)) \geq \omega_{ge}(K,L) - \sum_{i\in K}(v_i(x_i^*)\tilde{a}(x_i^*) - p_i)$$

Note that different choices of L lead to a different definitions of an ex-post core. However, we refer to any of these cores as *ex-post* cores, which reflects the fact that the core property is guaranteed for the point in time when availabilities have been realized.

Unfortunately, in our domain, the ex-post core can sometimes be empty.

Theorem 3. *In a domain with uncertainty, there does not exist a mechanism that is budget balanced and ex-post core-selecting.*

Proof. Consider a setting with two bidders $N = \{1, 2\}$ and a seller s. Assume that there are only two goods A and B. We assume that the first bidder has an value $v_1(\{A\})$ for a bundle $\{A\}$ and the second bidder has a value $v_2(\{AB\})$ for a bundle $\{AB\}$. We let $v_1(\{A\}) > 0$ and $v_2(\{AB\}) > 0$. We also let the costs of those goods to be equal to zero.

Given this specific set-up, there are seven different allocation rules we must consider: the allocation rule can allocate none of the items; it can allocate item A to bidder 1 and nothing to bidder 2; it can allocate item A to bidder 2 and nothing to bidder 1; etc. For each of these seven different allocation rules, there exists a joint probability mass function and a set of realized availabilities, such that we can construct an empty core.

We first consider the case where bidder 1 is allocated item A and bidder 2 is allocated the empty set. We let $a(\{A\})$ and $a(\{AB\})$ denote the availabilities of the bundles $\{A\}$ and $\{AB\}$ respectively, and we let f denote the corresponding joint probability mass function, which is defined as follows:

(a) $a(\{AB\}) = 0$

$a(A)$	$a(B) = 0$	$a(B) = 1$
0	γ_{000}	γ_{010}
1	γ_{100}	γ_{110}

(b) $a(\{AB\}) = 1$

$a(B) = 0$	$a(B) = 1$
γ_{001}	γ_{011}
γ_{101}	γ_{111}

Here $0 < \gamma_{ijk} < 1$, $i, j, k \in \{0, 1\}$, denotes the probability of the event that $a(\{A\}) = i$, $a(\{B\}) = j$, and $a(\{AB\}) = k$. Now we assume that $\tilde{a}(\{A\}) = 0$, $\tilde{a}(\{B\}) = 0$, $\tilde{a}(\{AB\}) = 1$. We can show that for any $L \subseteq 2^G$, the corresponding ex-post core is empty. Indeed, $\forall L \subseteq 2^G : \mathbb{E}[a(\{AB\})|\tilde{a}(S), S \in L] > 0$. Thus, the ex-post core constraint is then:

$$\omega_{ge}(\{s, 2\}, L) =$$
$$v_2(\{AB\}) \cdot \mathbb{E}[a(\{AB\})|\tilde{a}(S), S \in L] \leq 0 + p_s = p_1.$$

Here, p_1 and p_s are payments of the first bidder and the seller respectively (which are equal, given budget balance, i.e., $p_1 + p_s = 0$). Now, taking into account that $v_2(\{AB\}) > 0$ but $p_1 \leq v_1(\{A\})\tilde{a}(\{A\}) = 0$ (another core constraint), we get a contradiction. The proof for the other six possible allocations is analogous.

□

4.3 Framework for Execution-Contingent Mechanisms

Given that the ex-ante core can lead to large IR violations (as we will show in Section 5), and given the impossibility result regarding ex-post cores from Theorem 3, this raises the question which core to consider in practice. In this section, we put forward the idea of designing mechanisms that select payments inside an ex-post core *whenever this core is not empty*. Specifically, we design two *execution-contingent* payment rules. As we will later show in Section 5, for these payment rules, the empty core cases happen relatively rarely.

We will now first define a *mechanism framework* for execution-contingent payment rules which we then instantiate in two different ways:

Execution-Contingent Core-selecting Mechanism Framework

- *Framework parameters*:
  1. Reference point: p^*
  2. Core constraints: $Core^*$
- *Allocation rule*: select $x^* \in \text{argmax}_x \mathbb{E}_f[SW(x)]$
- *Payment rule*:

$$p = \begin{cases} p \in \text{argmin}_{p \in \Pi} ||p - p^*||_2 & \text{if Core}^* \cap \text{IR} \neq \emptyset \\ p* & \text{else} \end{cases}$$

where $\Pi = Core^* \cap IR \cap MRC$.

The first framework parameter, p^*, is a reference point, which we will either instantiate to $p^* = p^{\text{ECC-VCG}}$ or $p^{\text{ECR-VCG}}$, as defined by the ECC-VCG and ECR-VCG mechanisms. The second framework parameter is a set of core constraints $Core^*$, which we will accordingly instantiate to ECC-Core or ECR-Core, to be defined next. The ultimate core-selecting payment rule then first tries to find a payment vector p that minimizes the Euclidean distance to the reference point p^*, from among all payment vectors in the core that also minimize the revenue for the seller (the so-called *minimize-revenue constraint (MRC)*). We defined the mechanism framework this way to be analogous to the *Quadratic rule* [Day and Cramton, 2012], i.e., the core-selecting rule most commonly used in practice. In a domain with uncertain availabilities, however, the core (i.e., the intersection of the core constraints and the IR constraints) can be empty. In this case, we charge the reference point p^*, which will then be outside the core. Thus, all of these execution-contingent mechanisms are only ex-post core selecting whenever the core is non-empty. In Section 5, we will analyze how often such "empty core" cases occur in equilibrium.

4.4 ECC-Core Mechanism

Before we can introduce the ECC-Core mechanism, we need one more definition.

Definition 10. *An expected coalitional value $\omega_e(K)$ of a coalition $K \subset N$ is the maximum expected welfare the coalition can achieve knowing realized availabilities of allocated bundles. Formally,*

$$\omega_e(K) = \max_x \mathbb{E}_f \Big[\sum_{i \in K} (v_i(x_i) - c(x_i)) a(x_i) \Big| \tilde{a}(x_i^*), i \in W_{x^*} \Big]$$

Note, that this is a special case of Definition 9 when assuming $L = \{S \subseteq x_i^*, i \in W_{x^*}\}$[1].

Knowing realized availabilities of allocated bundles, bidders have a total opportunity cost of $\sum_{i \in K}(v_i(x_i^*)\tilde{a}(x_i^*) - p_i)$ for joining coalition K. If they decide to join, then in expectation they can achieve the total welfare of $\omega_e(K)$ and thus they can provide at most $\omega_e(K) - \sum_{i \in K}(v_i(x_i^*)\tilde{a}(x_i^*) - p_i)$ of additional value to the seller, which gives rise to the following set of core constraints:

ECC-Core Constraint $\forall K \subset N$:

$$\sum_{i \in W_{x^*}} (p_i - c(x_i^*)\tilde{a}(x_i^*)) \geq \omega_e(K) - \sum_{i \in K} (v_i(x_i^*)\tilde{a}(x_i^*) - p_i)$$

By plugging these constraints as $Core^*$ into the mechanism framework together with $p^* = p^{\text{ECC-VCG}}$ as the reference point, we obtain a full specification of the *execution-contingent conditional core (ECC-Core)* mechanism. The following example demonstrates how a coalition imposes a core constraint:

Example 2. *Consider the setting from Example 1, with coalition $K = \{b_3, s\}$. The expected coalitional value for this coalition is $\omega_e(K) = 0.3 \cdot 0.4 = 0.12$. The corresponding core constraint is $p_1 + p_2 \geq 0.12$.*

4.5 ECR-Core Mechanism

If we know the realizations of the availabilities of *all* bundles at the consumption time, and not only of those *allocated*, then we can use more accurate execution-contingent core constraints:

ECR-Core Constraint $\forall K \subset N$:

$$\sum_{i \in W_{x^*}} (p_i - c(x_i^*)\tilde{a}(x_i^*)) \geq \omega(K) - \sum_{i \in K} (v_i(x_i^*)\tilde{a}(x_i^*) - p_i),$$

where $\omega(K) = \max_x \sum_{i \in K}(v_i(x_i) - c(x_i))\tilde{a}(x_i)$

Note that this definition is a special case of Definition 9 assuming $L = \{S : S \in 2^G\}$. To get a full specification of the mechanism we use these core constraints together with $p^{\text{ECR-VCG}}$ as parameters for the execution-contingent mechanism framework. As we will show in Section 5 exploiting this additional knowledge can significantly decrease the rate of IR violations.

The following theorem shows that ECC-VCG and ECR-VCG provide lower bounds for the corresponding core-selecting payment rules. This is useful to know, because it also implies that using these payment vectors as reference points in the overall mechanism framework makes sense.

Theorem 4. *ECC-Core and ECR-Core payments are lower-bounded by ECC-VCG and ECR-VCG payments respectively.*

Proof. Consider ECC-Core mechanism. If the ECC-Core is empty, then the mechanism charges ECC-VCG payments which are trivially lower-bounded by ECC-VCG payments. If the ECC-Core is not empty, then consider a coalition $K = N \backslash \{k\}$, where $k \in W_{x^*}$. In this case

$$\mathbb{E}[\omega(K)] = \mathbb{E}_f[SW(x^{-i})|\tilde{a}(x_i^*), i \in W_{x^*}]$$

Then,

$$p_i \geq \mathbb{E}_f[SW(x^{-i})|\tilde{a}(x_i^*), i \in W_{x^*}] - \sum_{i \in W_{x^*} \backslash k} (v_i(x_i^*) - c(x_i^*))\tilde{a}(x_i^*) - c(x_k^*)\tilde{a}(x_k^*) = p_i^{\text{ECC-VCG}}.$$

The proof for ECR-Core is analogous. □

[1]Depending on the information structure of the domain, the availability that is revealed to the mechanism may only be $\tilde{a}(x_i^*), i \in W_{x^*}$ or $\tilde{a}(S), \forall S \subseteq x_i^*, i \in W_{x^*}$

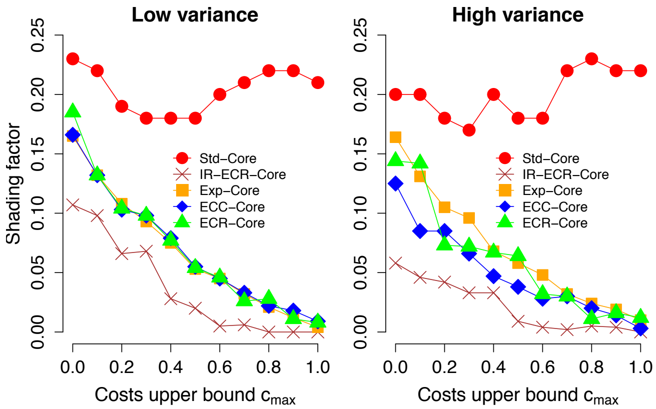

Figure 2: Additive shading factors of local bidders depending on costs distribution $c \sim U[0, c_{max}]$ for two levels of variance of the availabilities a.

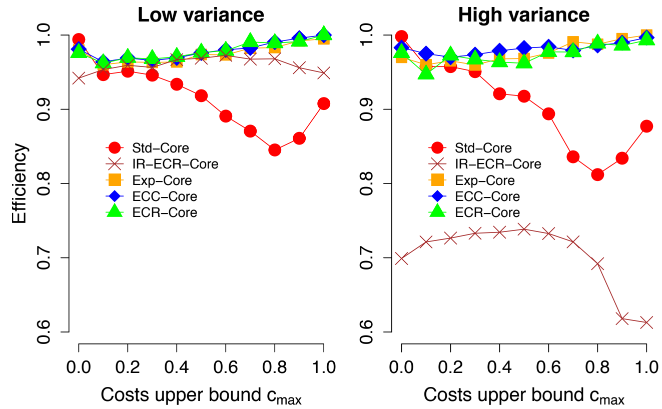

Figure 3: Efficiency depending on costs distribution $c \sim U[0, c_{max}]$ for two levels of variance of the availabilities a.

5 Comparison in Bayes-Nash Equilibrium

Our core-selecting payment rules are not strategyproof in expectation, and indeed, there do not even exist strategyproof core-selecting rules in domains without uncertainty [Goeree and Lien, 2016]. For this reason, we must analyze the properties (efficiency, IR violations, etc.) of our rules in equilibrium. The equilibrium concept we adopt for our auction domain is a *Bayes-Nash equilibrium*, i.e., where each bidder knows his own value, but only knows the distribution over other bidders' values, the seller's costs, and the availabilities.

Unfortunately, deriving the Bayes-Nash equilibrium of core-selecting payment rules analytically is very complex, and only feasible for very simple settings [Goeree and Lien, 2016]. For this reason we follow the approach by [Lubin and Parkes, 2009] and [Lubin *et al.*, 2016], and use a computational approach to find *approximate* BNEs for our rules. Concretely, we restrict the strategy space of the agents to *additive* shading strategies, and then use an algorithm based on fictitious play which, using an iterative best response method, converges to an ϵ-BNE in this restricted strategy space. Specifically, all the equilibria we report in this paper are ϵ-BNEs with ϵ=0.01.

5.1 Benchmark Rules

In addition to the ECC-Core and ECR-Core mechanisms, we also study the following three benchmark rules:

- **Std-Core Mechanism:** This refers to the standard core-selecting payment rule (the Quadratic rule as defined in [Day and Cramton, 2012]). The allocation rule selects an allocation assuming that all items are available, and payments are computed at allocation time.
- **Exp-Core Mechanism** This mechanism uses the ex-ante core as defined in Section 4.1, and then picks payments from this ex-ante core using the Quadratic rule [Day and Cramton, 2012]. Note that this mechanism satisfies individual rationality in expectation.
- **IR-ECR-Core Mechanism:** This refers to a rather artificial but still interesting mechanism. The allocation rule first maximizes expected social welfare, but then checks for every possible realization of the availabilities whether the resulting payments will satisfy ex-post IR. Only if this is the case will the allocation be made - otherwise no trade will happen. Thus, this rule is guaranteed to satisfy ex-post IR by design.

5.2 The Local-Local-Global (LLG) Domain

We study the well-known LLG domain. In this domain, there are two goods A and B. There are two local bidders who are each interested in one of the goods, and their values are drawn uniformly from $[0, 1]$, and there is a global bidders who is only interested in the bundle consisting of both goods, and his value for the whole bundle is drawn uniformly from $[0, 2]$. The seller's costs are drawn uniformly from $[0, c_{max}]$, where c_{max} is a parameter which we vary in our analysis.

It is easy to show that under Exp-Core, ECC-Core, ECR-Core, and IR-ECR-Core, if the global bidder gets allocated, he is charged the respective generalized version of the VCG payment. Thus, by Theorem 1, ECC-Core and ECR-Core are strategyproof in expectation for the global bidder, and the same result can be shown for Exp-Core and IR-ECR-Core. For this reason, we only need to compute the BNE strategies of the local bidders for those rules. For Std-Core, we also compute the BNE strategy for the global bidder.

5.3 Results

Strategies. Figure 2 shows the BNE strategies of the local bidders; on the left side for a domain with low variance in the availabilities of the goods, and on the right side for high variance in the availabilities of the goods. Note that these results are not "simulations," but that each point in these figures is the result of our BNE algorithm (and it takes about 8 hours to find a BNE for one payment rule on a machine with 20 cores).

We see that Std-Core has the worst incentives; the global bidder actually also shades (not shown in Figure 2), with a shading factor roughly twice as high as that of the local bidders. ECC-Core, ECR-Core and Exp-Core have very similar incentives, and IR-ECR-Core has the best incentives. Furthermore, the higher the costs, the lower the shading factors (except for Std-Core), which makes sense, because with higher costs, the opportunities for trade get smaller, and thus it gets more risky

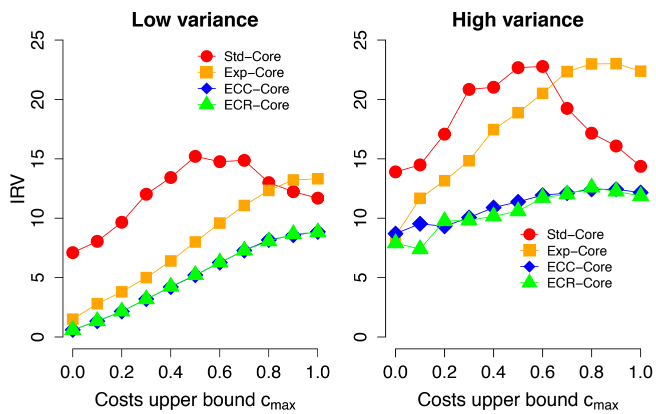

Figure 4: Rate of IR violations (%) depending on costs distribution $c \sim U[0, c_{max}]$ for two levels of variance of the availabilities a.

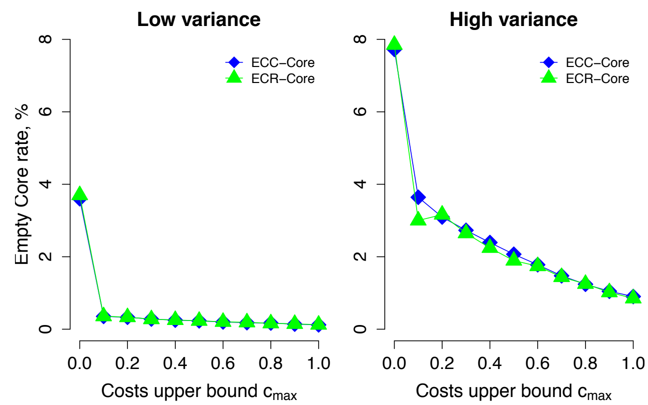

Figure 5: Rate empty core cases (%) depending on costs distribution $c \sim U[0, c_{max}]$ for two levels of variance of the availabilities a.

to shade. We also see that the shading factors are lower in the high variance domain, which makes sense, because the higher the uncertainty, the more risky it is to shade.

Efficiency. Figure 3 shows the efficiency achieved by all mechanisms (we simulate 1 million auctions using the compuated BNE strategies to calulate the efficiency). We see that our new mechanisms achieve higher efficiency than both, Std-Core as well as IR-ECR-Core.[2] By comparing the low variance and high variance domains, we see in what sense the IR-ECR-Core mechanism is really just a straw-man: while it works well in the low variance domain, in the high variance domain the mechanism has to cancel a lot of allocations because it cannot guarantee ex-post IR for all possible realizations of the availabilities, which drives down efficiency. In contrast, our execution-contingent mechanisms can handle this uncertainty.

Rate of Ex-post IR Violations. Figure 4 shows the rate of ex-post IR violations. Again, Std-Core performs worst. But this analysis now also demonstrates the advantages of ECC-Core and ECR-Core over Exp-Core. While all three mechanisms had good incentives and high efficiency, we now see that ECC-Core and ECR-Core have a significantly lower rate of ex-post IR violations, which makes sense because ECC-Core and ECR-Core are execution-contingent. This advantage is even more pronounced for the high-variance domain. In this domain, we even see a small advantage of ECR-Core over ECC-Core, which makes sense, because ECR-Core takes even more information into account when computing payments.

Empty Core Analysis. While in a domain without uncertainty, there always exists a price vector in the core, one of the interesting and perhaps unexpected features of our domain is that the ECC-Core as well as the ECR-Core can be empty. For this reason, we also study how often this happens in our domain. Figure 5 shows that in the low variance domain, the core is typically not empty, especially if the costs are non-zero then the probability of an empty core is close to 0. However, in the high variance domain, the rate of empty core cases is significantly higher. This is explained by the fact that the higher the variance in availabilities of bundles, the more flexible the core constraints are, and thus the more likely they can be in conflict with each other (and thus lead to an empty core).

[2]There is one exception: with zero costs, Std-Core achieve 99% efficiency while our rules achieve 98% efficiency. This happens because under our rules, the global bidder plays truthful and the local bidders shave, which leads to an efficiency loss. Under Std-Core, all bidders shave roughly proportionally, which leads to almost no efficiency loss in equilibrium (with zero costs). However, this peculiarity of the LLG domain does not generalize to larger domains.

6 Conclusion

In this paper, we have studied mechanisms for combinatorial auctions with uncertain availabilities of goods. We have introduced two execution-contingent core-selecting payment rules which both satisfy IR in expectation. Furthermore, we have performed an extensive computational Bayes-Nash equilibrium analysis, comparing our new rules with three benchmark rules to study the trade-off between different mechanism design objectives. Our results show that, compared to a standard core-selecting auction, our rules have significantly higher efficiency and lower ex-post IR violations. Furthermore, comparing our two execution-contingent mechanisms, we conclude that the more information about realized availabilities the mechanism has, the more of it should be exploited in the computation of the payments, as this leads to an lower rate of IR violations. Moreover, by comparing our rules to an ex-post IR rule, we have shown that by relaxing the strict ex-post IR constraint, we can gain a lot in efficiency. Thus, if a small rate of ex-post IR violations is acceptable, then we recommend using one of our new payment rules for a combinatorial auction domain with uncertain availabilities of goods.

Acknowledgments. We would like to thank Timo Mennle for insightful discussions and the anonymous reviewers for helpful comments. This research is supported by the SNSF (Swiss National Science Foundation) project #153598.

References

[Ausubel and Baranov, 2014] Lawrence Ausubel and Oleg Baranov. A practical guide to the combinatorial clock auction. Working paper, 2014.

[Ausubel and Baranov, 2016] Lawrence M. Ausubel and Oleg V. Baranov. Core-selecting auctions with incomplete information, 2016. Working paper.

[Ausubel and Milgrom, 2006] Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 57–95. MIT Press, 2006.

[Ceppi *et al.*, 2015] Sofia Ceppi, Nicola Gatti, and Enrico H. Gerding. Mechanism Design for the Federation of Service Providers. Working Paper, 2015.

[Clarke, 1971] E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 2:19–33, 1971.

[Cramton and Doyle, 2015] Peter Cramton and Linda Doyle. An open access wireless market supporting public safety, universal service, and competition. Technical report, 2015.

[Cramton, 2013] Peter Cramton. Spectrum auction design. *Review of Industrial Organization*, 42(2):030–190, March 2013.

[Day and Cramton, 2012] Robert Day and Peter Cramton. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research*, 60(3):588–603, 2012.

[Day and Milgrom, 2008] Robert Day and Paul Milgrom. Core-selecting Package Auctions. *International Journal of Game Theory*, 36(3):393–407, 2008.

[Day and Raghavan, 2007] Robert W. Day and S. Raghavan. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science*, 53(9):1389–1406, September 2007.

[FCC, 2016] https://www.fcc.gov/general/700-mhz-public-safety-spectrum-0#block-menu-block-4; last accessed on: February 2, 2016.

[Goeree and Lien, 2016] Jacob Goeree and Yuanchuan Lien. On the impossibility of core-selecting auctions. *Theoretical Economics*, 2016. Forthcoming.

[Goetzendorf *et al.*, 2015] Andor Goetzendorf, Martin Bichler, Pasha Shabalin, and Robert W. Day. Compact Bid Languages and Core Pricing in Large Multi-item Auctions. *Management Science*, 61(7):1684–1703, 2015.

[Groves, 1973] T. Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, 1973.

[Lubin and Parkes, 2009] Benjamin Lubin and David C. Parkes. Quantifying the strategyproofness of mechanisms via metrics on payoff distributions. In *Proceedings of the 17th National Conference on Artificial Intelligence*, pages 74–81, 2009.

[Lubin *et al.*, 2016] Benjamin Lubin, Benedikt Bünz, and Sven Seuken. New Core-Selecting Payment Rules with Better Fairness and Incentive Properties. Working Paper, 2016.

[Milgrom, 2007] Paul Milgrom. Package auctions and exchanges. *Econometrica*, 75(4):935–965, 2007.

[Moor *et al.*, 2015] Dmitry Moor, Tobias Grubenmann, Sven Seuken, and Abraham Bernstein. A Double Auction for Querying the Web of Data. In *Proceedings of the Third Conference on Auctions, Market Mechanisms and Their Applications (AMMA)*, Chicago, USA, August 2015.

[Porter *et al.*, 2008] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172(15):1783 – 1799, 2008.

[Ramchurn *et al.*, 2009] Sarvapali D. Ramchurn, Claudio Mezzetti, Andrea Giovannucci, Juan A. Rodriguez-Aguilar, Rajdeep K. Dash, and Nicholas R. Jennings. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial Intelligence Research*, 35(1):119–159, June 2009.

[Sandholm, 2013] T. Sandholm. Very-large-scale generalized combinatorial multi-attribute auctions: Lessons from conducting $60 billion of sourcing. In Alvin E. Roth Nir Vulkan and Zvika Neeman, editors, *The Handbook of Market Design*, chapter 1. Oxford University Press, 2013.

[Seuken *et al.*, 2015] Sven Seuken, Pertti Juhani Visuri, Johanna Katariina Visuri, and Dan Zagursky. Systems and methods for allocating and pricing alternative network access resources with reserve prices, July 2015, U.S. Patent 0189580.

[Vickrey, 1961] William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1):8–37, 1961.