Truthfulness and Stochastic Dominance with Monetary Transfers

MARTIN HOEFER, Max-Planck-Institut für Informatik
THOMAS KESSELHEIM, Max-Planck-Institut für Informatik
BERTHOLD VÖCKING, RWTH Aachen University

We consider truthfulness concepts for auctions with payments based on first- and second-order stochastic dominance. We assume bidders consider wealth in standard quasi-linear form as valuation minus payments. Additionally, they are sensitive to risk in the distribution of wealth stemming from randomized mechanisms. First- and second-order stochastic dominance are well-known to capture risk-sensitivity, and we apply these concepts to capture truth-telling incentives for bidders.

As our first main result, we provide a complete characterization of all social-choice functions over binary single-parameter domains that can be implemented by a mechanism that is truthful in first- and second-order stochastic dominance. We show that these are exactly the social-choice functions implementable by truthful-in-expectation mechanisms, and we provide a novel payment rule that guarantees stochastic dominance. As our second main result we extend the celebrated randomized meta-rounding approach for truthful-in-expectation mechanisms in packing domains. We design mechanisms that are truthful in first-order stochastic dominance by spending only a logarithmic factor in the approximation guarantee.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Non-numerical Algorithms and Problems

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Stochastic Dominance, Mechanism Design, Truthfulness, Combinatorial Auctions, Randomized Meta-Rounding

1. INTRODUCTION

Many recent advances in algorithmic mechanism design stem from the use of randomization in the design of social-choice rules. Especially in the prior-free setting, a number of deep and non-trivial techniques have been proposed to design mechanisms with optimal or near-optimal welfare guarantees for different variants of combinatorial auctions. A fundamental issue with randomization, however, is the definition of incentives over such lotteries. So far the attention in algorithmic mechanism design has focused on two notions of truthfulness in the presence of randomization.

— A mechanism is *truthful in expectation* if reporting the true valuation maximizes the expected quasi-linear utility for bidder i, i.e., the expected difference of outcome valuation

A preliminary version of this article appeared in the proceedings of the 14th ACM Conference of Electronic Commerce (EC 2013) [Hoefer et al. 2013].
This work is supported by DFG through Cluster of Excellence MMCI, grant Ho 3831/3-1, and UMIC Research Center at RWTH Aachen University, as well as by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD).
Author's addresses: M. Hoefer, Max-Planck-Institut für Informatik, Saarbrücken, Germany; T. Kesselheim, Max-Planck-Institut für Informatik, Saarbrücken, Germany; B. Vöcking, Department of Computer Science, RWTH Aachen University, Germany.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

v_i minus payment p_i. This ensures truthfulness if bidders care only about the *expectation* of the induced lottery over $v_i - p_i$.

— For systems with bidders that take more parameters of the distribution function into account, the standard concept is *universal truthfulness*. A mechanism is *universally truthful* if it is a probability distribution over truthful deterministic mechanisms. That is, even when a bidder knows all random decisions in advance, reporting the true valuation is a dominant strategy. The randomization is not used to ensure truthfulness but only to guarantee bounds on solution quality.

Universal truthfulness is quite a strong restriction over truthfulness in expectation as it requires a mechanism to generate incentives that work for any outcome of the random coin flips. Not surprisingly, there exist domains in which universally truthful mechanisms are significantly inferior to truthful-in-expectation mechanisms in terms of worst-case solution quality [Dobzinski and Dughmi 2013; Dughmi et al. 2011].

In contrast, truthfulness in expectation can only incentivize truth-telling for *risk-neutral* bidders. In this paper, we explore this issue by considering notions of randomized truthfulness based on stochastic dominance. In our approach, random coin flips are not revealed to the bidders in advance and can therefore be used to ensure truthful-telling, similar to truthfulness in expectation. However, we do not rely on the strong assumption of risk neutrality: Instead of comparing random variables only based on their expectation, we resort to stochastic dominance relations. Let us point out that we do not make any assumptions about the users attitude towards risk – which separates our approach from (recent) related work [Fu et al. 2013; Hu et al. 2010; Maskin and Riley 1984]. In particular, we investigate approaches that transform truthful in expectation mechanisms into mechanisms that are truthful with respect to stochastic dominance for users that might have any (unknown) attitude to risk.

In economic theory, stochastic dominance is a well-established concept to compare the returns and the riskiness of random outcomes [Mas-Colell et al. 1995, pp 194–199]. Random variable X is said to *first-order stochastically dominate* random variable Y if it yields "unambiguously higher return", formally $\mathbf{Pr}[X \geq a] \geq \mathbf{Pr}[Y \geq a]$ for all $a \in \mathbb{R}$. This is an indubitably convincing concept as any rational agent would prefer the first-order stochastic dominant distribution, independent of her perception of risk. In addition, we also consider the weaker notion of *second-order stochastic dominance*, where X second-order stochastically dominates Y if the integral over the distribution functions is dominated (for a formal definition see below). Intuitively, a risk-averse agent would prefer the distribution X against the second order dominated distribution Y, but a risk-seeking bidder might prefer Y.

The mechanisms that we design will be truthful in first-order stochastic dominance. That is, the random variable determined by the difference of valuation and payment that arises when telling the true type first-order stochastically dominates the one that arises when reporting a false type. This way, neither risk-averse nor risk-seeking bidders have an incentive to lie. When given impossibility results, we will always show that there is even no mechanism that is truthful in second-order stochastic dominance. This means, risk-averse bidders can potentially profit from lying.

Notation and Preliminaries. The problem in mechanism design is to pick an outcome from a publicly known set $\mathcal{A}$. There are n participants or *bidders*. Bidder i has a private *valuation function* or *type* $v_i \colon \mathcal{A} \to \mathbb{R}$, which comes from a publicly known type space $v_i \in V_i$. We denote the cartesian product of types by $v \in V = V_1 \times \ldots \times V_n$. A *mechanism* is a pair (f, p), where f is a *social-choice function* that maps the declared types to an outcome $f \colon V \to \mathcal{A}$. The vector $p = (p_1, \ldots, p_n)$ consists of payment functions, where $p_i \colon V \to \mathbb{R}$ assigns a payment for bidder i. The *quasi-linear utility* or *wealth* for bidder i is $u_i(v) = v_i(f(v)) - p_i(v)$.

We are interested in randomized mechanisms, which are probability distributions over deterministic mechanisms. That is, analogously to randomized algorithms they provide outcomes and payments according to probability distributions. In the algorithmic mechanism design literature, there are three standard notions of a *truthful mechanism.* A deterministic mechanism (f, p) is *(deterministically) truthful* if wealth is optimized for a truthful bid

$$v_i(f(v)) - p_i(v) \geq v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i}) \ ,$$

for all $v_i, v_i' \in V_i$ and all $v_{-i} \in V_{-i}$. A randomized mechanism (f, p) is *universally truthful* if it is a probability distribution over deterministically truthful mechanisms. Here truthfulness is independent of the outcome of random choices within f and p. Finally, a mechanism (f, p) is *truthful in expectation* if the *expected wealth* is optimized for a truthful bid:

$$\mathbf{E}\left[v_i(f(v)) - p_i(v)\right] \geq \mathbf{E}\left[v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i})\right]$$

for all $v_i, v_i' \in V_i$ and all $v_{-i} \in V_{-i}$. The expectation is taken over the randomization of the mechanism. Obviously, every deterministically truthful mechanism is universally truthful, and every universally truthful mechanism is truthful in expectation. In general, these inclusions are strict.

The obvious drawback of truthfulness in expectation is the assumption that bidders must care linearly about the expectation of their wealth. In contrast, universal truthfulness is very demanding, and impossibility results in prominent domains like multi-unit auctions have been derived. We here propose two different notions based on stochastic dominance that capture the idea of *risk-sensitivity.* A mechanism (f, p) is *truthful in first-order stochastic dominance* if, for every $x \in \mathbb{R}$, the probability to obtain at least a wealth of x is maximized for a truthful bid:

$$\mathbf{Pr}\left[v_i(f(v)) - p_i(v) \geq x\right] \geq \mathbf{Pr}\left[v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i}) \geq x\right] \ ,$$

for all $v_i, v_i' \in V_i, v_{-i} \in V_{-i}, x \in \mathbb{R}$. If a mechanism is truthful in first-order stochastic dominance, then for all users that have any (unknown) intrinsic monotone utility function of their wealth, their expected utility is maximized by truth-telling.

Risk-averse users are traditionally assumed to act according to the expectation of a monotone and concave function. To incentivize every such risk-averse user, only the cumulative distribution function needs to be minimized in terms of the lower integral [Mas-Colell et al. 1995]. A mechanism (f, p) is *truthful in second-order stochastic dominance* if, for every $x \in \mathbb{R}$, the following holds:

$$\int_{\substack{t \in V_i \\ t \leq x}} \mathbf{Pr}\left[v_i(f(v)) - p_i(v) \leq t\right] dt \leq \int_{\substack{t \in V_i \\ t \leq x}} \mathbf{Pr}\left[v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i}) \leq t\right] dt \ ,$$

for all $v_i, v_i' \in V_i, v_{-i} \in V_{-i}, x \in \mathbb{R}$. Every universally truthful mechanism is truthful in first-order stochastic dominance, which in turn is truthful in second-order stochastic dominance, which in turn is truthful in expectation. In general, these inclusions are strict.

We consider notions of *individual rationality.* An individually rational mechanism gives an incentive to participate in the mechanism, where non-participation yields a utility of 0. Naturally, these incentives correspond to the aspects of the utility distributions that are vital in the definition of truthfulness. A deterministic truthful mechanism is individually rational if $u_i(v_i, v_{-i}) \geq 0$ for his true type v_i and every $v_{-i} \in V_{-i}$. Similarly, a universally truthful mechanism is individually rational if all deterministic mechanisms in the support are individually rational. A mechanism that is truthful in first- or second-order stochastic dominance is individually rational if non-participation is stochastically dominated, i.e., $\mathbf{Pr}\left[u_i(v_i, v_{-i}) < 0\right] = 0$ for the true type v_i and every $v_{-i} \in V_{-i}$. Finally, a truthful-in-expectation mechanism is individually rational if $\mathbf{E}\left[u_i(v_i, v_{-i})\right] \geq 0$ for the true type v_i and every $v_{-i} \in V_{-i}$.

In addition to individual rationality, we assume the mechanisms to make *no positive transfers* in domains with non-negative valuations. Note that if computational tractability is no concern, individual rationality and no positive transfers can be easily achieved in a deterministic truthful mechanism using VCG.

Our Contributions. Our first main contribution is a characterization of social-choice functions that admit truthful-in-stochastic-dominance mechanisms for binary single-parameter domains. Interestingly, these are precisely the ones that admit truthful-in-expectation mechanisms. We show this equivalence by giving a black-box transformation turning any truthful-in-expectation mechanism into a mechanism that is truthful in first-order stochastic dominance by only changing the payment scheme. In particular, we turn payments into a lottery that, in addition to truthfulness in first-order stochastic dominance, guarantees non-positive transfers and a notion of individual rationality. It also yields the same expected social welfare and expected revenue as the truthful-in-expectation mechanism.

We show that such a black-box transformation in this generality cannot exist for much more general domains. In particular, even for fractional single-parameter domains or (binary) single-minded combinatorial auctions there are mechanisms that are truthful in expectation but that cannot be turned into truthful-in-second-order-stochastic-dominance mechanisms under standard assumptions without changing the social-choice function.

A similar problem arises when considering truthful-in-expectation mechanisms derived by the meta-rounding framework by Lavi and Swamy [2011]. This is due to the fact that optimal LP solutions lack certain monotonicity conditions required for truthful-in-stochastic-dominance mechanisms. Our second main contribution is an alternative framework that also uses meta-rounding to derive mechanisms that are truthful-in-stochastic-dominance. In contrast to the standard approach, we do not use the optimal LP solution but an approximate one. We lose only a factor of $O(\log m + \log n)$ in the approximation guarantee of social welfare, where n is number of bidders and m is a parameter of the problem, e.g., the number of items in a (multi-unit) combinatorial auction.

Related Work. First-order stochastic dominance is a classic truthfulness property in the social-choice literature, as it implies truthfulness in expectation for every realization of ordinal preferences in cardinal utility values. Gibbard [1977] showed an impossibility theorem for randomized mechanisms without payments. The random dictatorship mechanism is the unique social-choice rule which is truthful in first-order stochastic dominance and never puts positive probability on Pareto-dominated solutions. More recently, truthfulness in first-order stochastic dominance for randomized social choice has been considered, e.g., by Aziz et al. [2013] who characterize voting rules with respect to several less demanding concepts of efficiency and truthfulness, including weaker versions of truthfulness in first-order stochastic dominance. For another approach to preferences over distributions generalizing stochastic dominance see Peters et al. [2010].

Over the last decade, there has been increased interest in characterizing randomized social-choice mechanisms for restricted domains of allocation problems. Initial contributions that consider truthfulness in first-order stochastic dominance discuss the random priority mechanism, a repeated version of random dictatorship [Zhou 1990; Abdulkadiroglu and Sönmez 1998]. Bogomolnaia and Moulin [2001] and Katta and Sethuraman [2006] introduce and examine a probabilistic serial mechanism that satisfies stronger efficiency guarantees but a weaker variant of truthfulness in stochastic dominance. In contrast to the classic approach where the distribution obtained by truth telling should weakly dominate all distributions obtained from lying, the weaker variant only requires that the truthful distribution is never strictly dominated by any distribution obtained from lying. For similar randomized mechanisms in allocation with dichotomous preferences see Bogomolnaia and Moulin [2004]. More recently, a variety of works are extending this line of research. In addition, single-peaked preferences have become a domain of interest as, notably, Ehlers et al. [2002] extended the characterization of Moulin [1980] to randomized mechanisms that are truthful in first-order stochastic dominance.

First- and second-order stochastic dominance are well-known to capture the incentives for risk-sensitive agents in markets. For a basic treatment of fundamental connections see, e.g., [Mas-Colell et al. 1995, pp 194–199]. Risk aversion is a recent trend in algorithmic mechanism design [Dughmi and Peres 2012; Fu et al. 2013; Bhalgat et al. 2012; Sundararajan and Yan 2010]. Most closely related to our work is a note by Dughmi and Peres [2012], who consider risk-averse bidders that optimize expectation over *a concave utility function.* They show that every social-choice rule of a truthful-in-expectation mechanism for risk-neutral bidders can be turned into a truthful-in-expectation mechanism for risk-averse bidders by designing appropriate payments. In our terms, there is a black-box transformation for every truthful-in-expectation mechanism into one that is truthful in second-order stochastic dominance. While being very general, the transformation does not guarantee payments to flow only in one direction. In fact, we are able to show a different transformation in single-parameter domains, where we even obtain first-order stochastic dominance. Additionally, we can guarantee a natural individual rationality constraint and non-negativity of payments. In contrast, we give an example for single-minded combinatorial auctions where the result of [Dughmi and Peres 2012] is impossible to achieve in combination with these conditions.

In the mechanism design literature, risk aversion has often been considered for single-item auctions. In this context, several works have characterized revenue and differences between first- and second-price auctions [Maskin and Riley 1984; Esö and Futo 1999; Hu et al. 2010]. Recently, Fu et al. [2013] consider a Bayesian setting and design prior-independent auctions for revenue maximization with risk averse bidders. Valuation of an item is drawn from a prior distribution, and risk aversion takes a very particular form, i.e., the concave utility function grows linearly up to a capacity bound and stays constant afterwards. Given knowledge about the capacity, it is possible to construct optimal auctions and prior-independent constant-factor approximations of it. These auctions outperform the optimal revenue in the risk-neutral case significantly by exploiting knowledge about the form of risk averseness. In contrast, similar to Dughmi and Peres [2012] we assume that the bidders attitude towards risk is unknown, so our auctions cannot exploit specific knowledge about risk sensitivity. In particular, risk neutrality of bidders remains a possible option, and hence our mechanisms must fulfill strictly stronger conditions than truthful-in-expectation mechanisms.

On the technical side, some of our proofs are related to recent advances in algorithmic aspects of combinatorial auctions. We adjust the celebrated randomized meta-rounding technique of Lavi and Swamy [2011] and combine it with a fractional overselling algorithm recently studied by Krysta and Vöcking [2012]. Our truthfulness concepts lie in between truthfulness in expectation and universal truthfulness, for which separation and characterization results in terms of complexity and approximation have recently attracted significant interest. For a discussion of the recent literature see, e.g. [Vöcking 2012; Dughmi 2012], and for connections to risk aversion see [Dughmi and Peres 2012].

Outline. The black-box transformation for single-parameter domains is presented in Section 2. The meta-rounding approach for multi-parameter domains is presented in Section 3. In Section 4 we present a number of applications of our technique. In Section 5 we discuss our approach in the context of related work and conclude in Section 6 with interesting avenues for future work.

2. SINGLE-PARAMETER AND SINGLE-MINDED DOMAINS

Single-Parameter Domains. For binary single-parameter domains, each valuation can be represented by a private scalar $v_i \in \mathbb{R}$. For each bidder i, there is a publicly known interval such that $v_i \in [v_i^0, v_i^1]$. In addition, for each bidder there is publicly known subset $W_i \subset \mathcal{A}$ labeled as *winning outcomes.* For each winning outcome, bidder i's valuation is v_i, for every other outcome it is 0. For a given randomized mechanism, we denote by $w_i(v)$ the probability that it outputs an outcome from W_i. The following classic characterization of truthful-in-expectation mechanisms of Myerson [1981] is taken from Nisan [2007]. For simplicity, we consider *normalized* mechanisms, for which $w_i(v_i^0, v_{-i}) = 0$ and $p_i(v_i^0, v_{-i}) = 0$.

THEOREM 2.1 ([MYERSON 1981; NISAN 2007]). *A normalized mechanism (f, p) over a binary single parameter domain is truthful in expectation if and only if for every bidder i and every v_{-i} we have*

(1) *the winning probability $w_i(v_i, v_{-i})$ is monotonically non-decreasing in v_i, and*
(2) $\mathbf{E}\left[p_i(v_i, v_{-i})\right] = v_i \cdot w_i(v_i, v_{-i}) - \int_{v_i^0}^{v_i} w(t, v_{-i}) dt.$

Our main result in this section shows that we can strengthen the "if" part of this characterization as follows. Given any randomized social-choice function f over a binary single parameter with monotonically non-decreasing winning probabilities, we can devise randomized payments p so that (f, p) is truthful in first-order stochastic dominance.

THEOREM 2.2. *A normalized mechanism (f, p) over a binary single parameter domain is truthful in first-order stochastic dominance if for every bidder i and every v_{-i} we have*

(1) *the winning probability $w_i(v_i, v_{-i})$ is monotonically non-decreasing in v_i, and*
(2) $p_i(v_i, v_{-i}) = 0$ *if* $f(v_i, v_{-i}) \notin W_i$*; otherwise* $p_i(v_i, v_{-i})$ *is drawn at random by setting* $p_i(v_i, v_{-i}) = \min\{p \mid w_i(p, v_{-i}) \geq \beta\}$*, where* β *is drawn uniformly at random from* $[0, w_i(v_i, v_{-i})]$.

Note that monotonicity is required even for truthfulness in expectation. Therefore, this theorem implies that for every normalized mechanism (f, p) over a single-parameter domain that is truthful in expectation, there is an equivalent one (f, p') that is truthful in first-order stochastic dominance. The equivalent mechanism uses the same social-choice function f but possibly different, randomized payment rules.

PROOF OF THEOREM 2.2. The intuition behind this payment rule is that a bidder can barely influence the probability distribution of her payment. For example, by reporting a higher valuation, the overall probability of winning and being charged a low payment does not change. Only probability mass is shifted from non-winning outcomes to winning outcomes with high payments. As we will show, reporting the true valuation balances probability masses in such a way that it maximizes the probability of positive utility.

To show truthfulness in first-order stochastic dominance, we have to show that for all $x \in \mathbb{R}$, we have

$$\mathbf{Pr}\left[u_i(v_i, v_{-i}) \geq x\right] \geq \mathbf{Pr}\left[u_i(v_i', v_{-i}) \geq x\right] \quad .$$

As v_{-i} remains fixed throughout this proof, we omit this argument from now on to keep notation simple.

Let us first consider the case that $x \leq 0$. Observe that the payment is upper bounded by the bid, so the utility will always be non-negative when bidding truthfully. Formally, this means $\mathbf{Pr}\left[u_i(v_i) < 0\right] = 0$, or in other words $\mathbf{Pr}\left[u_i(v_i) \geq x\right] = 1$ for all $x \leq 0$. Therefore, we trivially have $\mathbf{Pr}\left[u_i(v_i) \geq x\right] \geq \mathbf{Pr}\left[u_i(v_i') \geq x\right]$ for all $v_i' \in \mathbb{R}$ for $x \leq 0$.

In case $x > 0$, two events have to occur to have $u_i(v_i') \geq x$. First, bidder i needs to be assigned a winning outcome, i.e., $f(v_i') \in W_i$. The probability for this event is $w_i(v_i')$. Second, the payment has to be at most $y = v_i - x$. Let us consider this probability given that $f(v_i') \in W_i$. The payment will never exceed v_i', so we know that $\mathbf{Pr}\left[p_i(v_i') \leq y \mid f(v') \in W_i\right] = 1$ for $y \geq v'$. For $y \leq v'$, in turn, the payment is bounded by y if and only if some β is chosen such that $w_i(y) \geq \beta$. The probability for this event is $\frac{w_i(y)}{w_i(v_i')}$. That is, in total, we get

$$\mathbf{Pr}\left[u_i(v_i') \geq x \mid f(v_i', v_{-i}) \in W_i\right] = \begin{cases} 1 & \text{if } v_i - x \geq v_i' \\ \frac{w_i(v_i - x)}{w_i(v_i')} & \text{if } v_i - x \leq v_i' \end{cases}$$

and therefore

$$\begin{aligned}\mathbf{Pr}\left[u_i(v_i') \geq x\right] &= \begin{cases} w_i(v_i') & \text{if } v_i - x \geq v_i' \\ w_i(v_i - x) & \text{if } v_i - x \leq v_i' \end{cases} \\ &= \min\{w_i(v_i'), w_i(v_i - x)\} \ .\end{aligned}$$

In particular, as w_i is non decreasing, we have $\mathbf{Pr}\left[u_i(v_i) \geq x\right] = w_i(v_i - x) \geq \min\{w_i(v_i'), w_i(v_i - x)\} = \mathbf{Pr}\left[u_i(v_i') \geq x\right]$ for all $x > 0$ and all $v_i' \in \mathbb{R}$. This proves the claims. □

Observe that our mechanisms are individually rational because no bidder is charged more than his bid. Similarly, by adjusting the payments we can also guarantee individual rationality for the corresponding truthful-in-expectation mechanisms (e.g., by dividing the payment of Theorem 2.1 by $w_i(v_i, v_{-i})$ and assigning it only in case $f(v) \in W_i$). Finally, the mechanisms also have no positive transfers for $v_0^i \geq 0$ because the minimum payment is bounded by v_0^i.

Single-Minded Combinatorial Auctions. When we consider the slightly more general setting of single-minded combinatorial auctions, we fail to obtain a similar general correspondence result as for single-parameter domains. In this scenario we have n bidders and m indivisible items. Each bidder i has a private value $v_i \in \mathbb{R}$ and a private set $S_i \subseteq [m]$. An outcome is an allocation of items to bidders. If bidder i gets assigned a set $S \supseteq S_i$ of items, his valuation for the outcome is v_i, otherwise it is 0. The main difference to single-parameter domains is that in single-minded combinatorial auctions the sets S_i are also private knowledge.

The following characterization of mechanisms that are truthful in expectation is easy to derive; it collapses to the well-known monotonicity criterion for deterministic mechanisms [Nisan 2007]. We omit the straightforward proof.

THEOREM 2.3. *A normalized mechanism (f, p) for single-minded combinatorial auctions is truthful in expectation if and only if for every bidder i with bid (v_i, S_i) and every v_{-i}, S_{-i} we have*

(1) $w_i(v_i, S_i, v_{-i}, S_{-i}) \geq w_i(v_i', S_i, v_{-i}, S_{-i})$ *for* $v_i' \leq v_i$,

(2) $\int_0^{v_i} w_i(t, S_i, v_{-i}, S_{-i}) dt \geq \int_0^{v_i} w_i(t, S_i', v_{-i}, S_{-i}) dt$ *for* $S_i' \supseteq S_i$,

(3) $p_i(v_i, S_i, v_{-i}, S_{-i}) = v_i \cdot w_i(v_i, S_i, v_{-i}, S_{-i}) - \int_{v_i^0}^{v_i} w(t, S_i, v_{-i}, S_{-i}) dt$

The characterization requires that the *integral over* the winning probability is monotonically decreasing in the declared set. As a direct corollary from our previous result we show that if the winning probability is monotone in both valuation and set, we can apply our previous technique and design payments for a mechanism that is truthful in stochastic dominance.

COROLLARY 2.4. *A normalized mechanism (f, p) for single-minded combinatorial auctions is truthful in first-order stochastic dominance if for every bidder i with bid (v_i, S_i) and every v_{-i}, S_{-i} we have*

(1) $w_i(v_i, S_i, v_{-i}, S_{-i}) \geq w_i(v_i', S_i, v_{-i}, S_{-i})$ *for* $v_i' \leq v_i$,

(2) $w_i(v_i, S_i, v_{-i}, S_{-i}) \geq w_i(v_i, S'_i, v_{-i}, S_{-i})$ *for* $S'_i \supseteq S_i$,
(3) $p_i(v_i, S_i, v_{-i}, S_{-i}) = 0$ *if* $f(v_i, S_i, v_{-i}, S_{-i}) \notin W_i$; *otherwise determine* $p_i(v_i, S_i, v_{-i}, S_{-i})$ *at random by setting* $p_i(v_i, S_i, v_{-i}, S_{-i}) = \min\{p \mid w_i(p, S_i, v_{-i}, S_{-i}) \geq \beta\}$, *where* β *is drawn uniformly at random from* $[0, w_i(v_i, S_i, v_{-i}, S_{-i})]$.

PROOF. Similarly to the proof of Theorem 2.2 we can argue that the payments for the true bid yield $\mathbf{Pr}[u(v,S) \geq x] = \min\{w(v,S), w(v_i - x, S_i, v_{-i}, S_{-i})\}$. Hence, a bidder i that keeps the true v_i and deviates to a larger set suffers because w is monotone in S_i. Keeping the true S_i and deviating to different valuations is stochastically dominated because of previous arguments for the single-parameter case. In combination, this shows the corollary. □

All truthful-in-expectation mechanisms can again be made individually rational. In addition, the mechanisms described in Corollary 2.4 are individually rational because payments do not exceed bids. However, note that Condition (2) in Theorem 2.3 only requires the integral of w to be monotone whereas Condition (2) in Corollary 2.4 requires monotonicity of the function w itself. Indeed, this brings about a fundamental difference to the case of binary single-parameter settings. In contrast to Theorem 2.2, there are monotone functions w for single-minded combinatorial auctions that admit implementation as truthful-in-expectation mechanisms but not as mechanisms that are individually rational and truthful in second-order stochastic dominance.

LEMMA 2.5. *There are single-minded combinatorial auctions with randomized social-choice functions* f *such that the following holds. There are payments* p *such that* (f, p) *is a normalized and individually rational mechanism which is truthful in expectation, but there are no payments* p' *such that* (f, p') *is a normalized and individually rational mechanism with no positive transfers that is truthful in second-order stochastic dominance.*

PROOF. We provide a simple counterexample to highlight the argument. There is one bidder and two items. For simplicity, we restrict all valuations to come from the interval $[0, 1]$. The winning probability based on f is such that

$$w(x, S) = \begin{cases} x & \text{if } S = \{1\} \\ 0 & \text{if } S = \{2\} \text{ or } S = \{1, 2\}, \text{ and } x < 0.5 \\ 1 & \text{otherwise.} \end{cases}$$

Note that w fulfills monotonicity in the integrals, because for all $x \in [0, 1]$

$$\int_{t=0}^{x} w(t, \{1\}) dt = \frac{x^2}{2} \quad \geq \quad \max\{0, x - 0.5\} = \int_{t=0}^{x} w(t, \{1, 2\}) dt \ .$$

Thus, we can turn this function into a mechanism that is truthful in expectation.

Now assume there is a mechanism (f, p') which is normalized, individually, with no positive transfers, and truthful in second-order stochastic dominance. Consider the case that the bidder has true type $(v, S) = (0.9, \{1\})$ and suppose he lies $(v', S') = (0.5, \{1, 2\})$. For individual rationality we require the payment to be at most the bid (for second-order one might exclude measure-zero sets of bids). Hence, for infinitely many bids in, say, $[0.5, 0.6]$ the bidder is always charged at most the bid. W.l.o.g. we assume this is true for declaring $(0.5, \{1, 2\})$, where the bidder is charged at most 0.5. This results in $\mathbf{Pr}[v(f(0.5, \{1, 2\})) - p'(0.5, \{1, 2\}) \geq 0.4] = 1$. In contrast, when being truthful, the probability $\mathbf{Pr}[v(f(0.9, \{1\})) - p'(0.9, \{1\}) \geq x] \leq 0.9$ for every $x \in (0, 0.9]$.

This means

$$\int_{t=0}^{0.4} \mathbf{Pr}[v(f(v, S)) - p'(v, S) \leq t] \, dt \geq 0.4 \cdot (1 - 0.9) \ ,$$

but

$$\int_{t=0}^{0.4} \mathbf{Pr}\left[v(f(v', S')) - p'(v', S') \leq t\right] dt = 0 \ \ ,$$

and contradicts second-order stochastic dominance. □

Observe that individual rationality requires payments to be upper bounded by bids (for second-order except measure-zero sets of bids), and our proof critically relies on this property. In fact, if we drop the conditions on the payments, the reduction by Dughmi and Peres [2012] implies a black-box transformation even for all general multi-parameter domains.

Fractional Single-Parameter Domains. Even if we restrict attention to single-parameter domains but drop the integrality condition, a similar black-box transformation as for binary domains becomes impossible. In a fractional single-parameter domain, each bidder i has a private value $v_i \in \mathbb{R}$ as above. Instead of subsets of winning outcomes W_i that yield value v_i, every outcome is now interpreted as a fractional winning assignment. For a vector b of bids, we obtain $f(b) = (x_1, \ldots, x_n)$, and the utility of bidder i becomes $u_i(b) = x_i \cdot v_i - p_i(b)$. Whereas we here consider randomized mechanisms, Archer and Tardos [2001] give a characterization of deterministic mechanisms that are truthful on this domain which is similar to the one by Myerson [1981].

Our result is similar to the previous one, and it also relies on individual rationality.

LEMMA 2.6. *There are fractional single-parameter domains with randomized social-choice functions f such that the following holds. There are payments p such that (f, p) is an individually rational mechanism which is truthful in expectation, but there are no payments p' such that (f, p') is an individually rational mechanism with no positive transfers that is truthful in second-order stochastic dominance.*

PROOF. We construct a similar counterexample as for combinatorial auctions. There is one bidder, and we restrict all valuations to come from the interval $[0, 1]$. The randomized function f yields a probability distribution over fractional outcomes for the bidder as follows

$$f(b_1) = \begin{cases} 0.1 & \text{with prob. 1, if } b_1 \leq 0.5, \\ 1 & \text{with prob. 0.5, if } b_1 > 0.5, \\ 0 & \text{with prob. 0.5, if } b_1 > 0.5. \end{cases}$$

The following payments yield a truthful-in-expectation mechanism.

$$p(b_1) = \begin{cases} 0 & \text{if } b_1 \leq 0.5, \\ 0.2 & \text{otherwise.} \end{cases}$$

Reporting a value below 0.5 gives expected utility $0.1v_1$. Reporting above 0.5 gives expected utility $0.5v_1 - 0.2$. Observe

$$0.1v_1 \ \geq \ 0.5v_1 - 0.2 \quad \Leftrightarrow \quad v_1 \leq 0.5 \ \ ,$$

and, thus, truthful reporting gives maximum expected utility.

Now assume there is a mechanism (f, p') which is normalized, individually, with no positive transfers, and truthful in second-order stochastic dominance. Consider the case that the bidder has true type $v = 0.7$ and suppose he lies $v' = 0$. For individual rationality we require the payment to be at most the bid (for second-order one might exclude measure-zero sets of bids). Hence, for infinitely many bids in, say, $[0, 0.01]$ the bidder is always charged at most the bid. W.l.o.g. we assume this is true for declaring 0, where the bidder is charged nothing. This results in $\mathbf{Pr}\left[v(f(0)) - p'(0) \geq 0.07\right] = \mathbf{Pr}\left[v(f(0)) - p'(0) \geq 0.7 \cdot 0.1\right] = 1$. In contrast, when being truthful, the probability $\mathbf{Pr}\left[v(f(0.7)) - p'(0.7) \geq 0.07\right] \leq 0.5$.

This means

$$\int_{t=0}^{0.05} \mathbf{Pr}\left[v(f(v)) - p'(v) \leq t\right] dt \geq 0.05 \cdot (1 - 0.5) \enspace ,$$

but

$$\int_{t=0}^{0.05} \mathbf{Pr}\left[v(f(v')) - p'(v') \leq t\right] dt = 0 \enspace ,$$

and contradicts second-order stochastic dominance. □

3. MULTI-PARAMETER PACKING DOMAINS

In the previous section we observed that even for single-minded combinatorial auctions there exist social-choice functions of truthful-in-expectation mechanisms that do not admit individually rational mechanisms that are truthful in second-order stochastic-dominance. Moreover, it can be shown that this holds even for social-choice functions derived via the randomized meta-rounding framework of Lavi and Swamy [2011]. In this section, we present an alternative approach to design mechanisms that are individually rational and truthful in first-order stochastic dominance under the same prerequisites as in [Lavi and Swamy 2011].

The general idea of the meta-rounding approach is as follows. One assumes that the underlying social-choice problem is given by maximizing social welfare via a packing integer program. Furthermore, one assumes that there is an efficient algorithm $\mathcal{A}$ verifying an integrality gap of $\alpha \leq 1$ in the natural LP relaxation. That is, given an optimal fractional solution x^* of social welfare $v^T x^*$ to the LP relaxation, $\mathcal{A}$ delivers a feasible integral solution for which social welfare is at least $\alpha \cdot v^T x^*$. This algorithm $\mathcal{A}$ is then used to decompose the optimal (fractional) LP solution x^* into polynomially many integral solutions as $x^* = \alpha \sum_\ell \lambda_\ell x^\ell$. To determine the output allocation, the weights λ_ℓ are interpreted as a probability distribution. By using scaled VCG payments with respect to the optimal fractional solution, one obtains a truthful-in-expectation mechanism.

The general idea in our approach remains based on randomized meta-rounding. In contrast, we do not decompose a scaled variant of x^*, our mechanisms compute and decompose a different solution x. It might be suboptimal but fulfills some important monotonicity conditions that allow to implement truthfulness in stochastic dominance.

More formally, we keep intuition and notation similar to combinatorial auctions introduced in the previous chapter. For each bidder i, there is a set $\mathcal{S}_i$ containing multiple entities $S \in \mathcal{S}_i$, one of which can be assigned to this bidder. In the integer program, variable $x_{i,S}$ captures if bidder i is assigned to entity S. Each bidder can be assigned at most one entity. In addition, there are packing constraints that restrict the possible entity-bidder combinations, i.e., the overall allocation. The goal is find an allocation that maximizes social welfare as given by the sum of valuations of assigned entities.

The underlying allocation problem can be stated by the following integer linear program:

$$\begin{array}{lrcll} \text{Max.} & v^T x & & & \\ \text{s.t.} & Ax & \leq & \mathbf{1} & \\ & \sum_{S \in \mathcal{S}_i} x_{i,S} & \leq & 1 & \text{for all } i \in [n], \\ & x_{i,S} & \in & \{0,1\} & \text{for all } i \in [n],\, S \in \mathcal{S}_i \enspace . \end{array}$$

The extension over ordinary combinatorial auctions lies in matrix A. We denote by m the number of rows in A. In combinatorial auctions each entity is a set of items, each row in A corresponds to one item, and the row constrains the item to be given to at most one bidder. We here assume A may contain arbitrary non-negative entries. Thereby it allows to express more complex constraints on the allowed combinations of entities that can be allocated to bidders.

ALGORITHM 1: Fractional Overselling MPU Algorithm

Initialize $x_{i,S} = 0$ for all i, S;
Add each bidder uniformly at random to set STAT $\subseteq [n]$ with probability $\frac{1}{2}$;
if STAT $\neq \emptyset$ **then** let $L := \max_{i \in \text{STAT}} \max_S v_{i,S}$ **else** let $L := 0$;
Set $p_0 = \frac{L}{4m}$;
foreach *constraint* $e \in [m]$ **do** $p_e^1 := p_0$;
foreach *bidder* $i \in [n] \setminus$ STAT **do**
  Set $S_i := D_i(p^i)$;
  Set $x_{i,S_i} = 1$;
  foreach *constraint* $e \in [m]$ **do**
    if $\sum_{i'} a_{e,(i',S_{i'})} \geq \log(4mb) + 3$ **then** set $p_e^{i+1} = \infty$;
    else update $p_e^{i+1} := p_e^i \cdot 2^{a_{e,(i,S_i)}}$;

Note that the number of columns in A equals the dimension of x and amounts to $\sum_{i\in[n]} |\mathcal{S}_i|$. We set b to the reciprocal minimal positive entry in the matrix A, but to at most n. Formally, this means $b = \min\left\{n, \max_{i,S,e;\ a_{e,(i,S)}\neq 0} \frac{1}{a_{e,(i,S)}}\right\}$. In our algorithms we use a scaling parameter $\gamma = \log(4mb) + 4$. Finally, let x^* be the optimal fractional solution to the integer program shown above, respectively. An integral solution computed by our mechanism is denoted by x_{ALG}.

We assume that we can access each bidder's valuation by a *demand oracle* $D_i(p)$, which returns the preferred entity given a vector of prices p. Given a vector of prices $(p_e)_{e\in[m]}$, it returns the $S \in \mathcal{S}_i$ maximizing $v_{i,S} - \sum_{e\in[m]} a_{e,(i,S)} p_e$. It may occur that this quantity is negative for all entities S. In this case, the bidder would prefer to opt out of the mechanism. To simplify notation, we add a virtual entity $\emptyset$ to each set $\mathcal{S}_i$ with $v_{i,\emptyset} = 0$ and all-zero coefficients in A. This way, no bidder will ever receive any set S for which the sum of prices exceeds his valuation.

In order to obtain the LP solution x which shall be decomposed, our mechanism uses a fractional overselling technique specified in Algorithm 1. It is similar to the one used by Krysta and Vöcking [2012]. The overall idea is to randomly divide the set of bidders into two sets and to sequentially make posted-price offers to the bidders in one set. The prices for outcomes are determined as follows: At each time, each of the constraints e in the matrix A has a different price. The price of outcome S depends on the values of $a_{(i,S),e}$, i.e., it depends on how allocating this outcome influences the constraints. The bidder then chooses his most preferred entity based on the prices using his demand oracle. Afterwards, the algorithm updates the prices for the next bidder. The other set of bidders, who do not participate in these auctions, are used to determine the initial prices.

LEMMA 3.1. *Algorithm 1 computes a vector x such that (1) $A \cdot x \leq \gamma \cdot \mathbf{1}$, (2) for each bidder i there is at most one $S_i \in \mathcal{S}_i$ such that $x_{i,S_i} = 1$ and all other entries are zero, and (3) $\mathbf{E}\left[v^T x\right] \geq \frac{1}{64} v^T x^*$.*

Parts (1) and (2), which are statements on feasibility of the solution, follow directly from the construction of the algorithm. To show part (3) we adapt the proof in [Krysta and Vöcking 2012] below. First, however, let us present our final mechanism given in Algorithm 2. It runs the overselling algorithm as a subroutine, which returns allocation x. Using part (1) of Lemma 3.1, $x' = {}^1\!/_\gamma \cdot x$ can be interpreted as a feasible fractional LP solution. Therefore, this solution x' can be treated like the optimal LP solution in the framework by Lavi and Swamy [2011] and be decomposed into polynomially many integral solutions. As payments we use the prices that were originally paid in the overselling algorithm.

ALGORITHM 2: Randomized Meta-Rounding Mechanism

Compute overselling solution (x, p) by running Algorithm 1;
Set $x' := \frac{1}{\gamma} \cdot x$, where $\gamma = \log(4mb) + 4$;
Decompose $x' = \alpha \sum_\ell \lambda_\ell x^\ell$ with $\sum_\ell \lambda_\ell = 1$;
Choose one of the x^ℓ at random with probability λ_ℓ each;
Assign entities according to chosen x^ℓ;
foreach *bidder* $i = 1, 2, \ldots, n$ **do**
  if $x^\ell_{i,S_i} = 1$ **then** charge bidder i payment $\sum_{e \in [m]} a_{e,(i,S_i)} p^i_e$;
  else charge bidder i no payment;

Let us bound the solution quality of Algorithm 2 under the assumption that Lemma 3.1 holds.

THEOREM 3.2. $\mathbf{E}\left[v^T x_{\mathrm{ALG}}\right] \geq \dfrac{\alpha}{64\gamma} \cdot v^T x^* = \dfrac{\alpha}{O(\log(m) + \log(n))} \cdot v^T x^*$.

PROOF. We have $v^T x' = \frac{1}{\gamma} v^T x$ and $\sum_\ell \lambda_\ell x^\ell = \alpha x'$. The expected social welfare of the computed solution is given by $v^T \sum_\ell \lambda_\ell x^\ell$. Assembling these bounds, we get $\mathbf{E}\left[v^T x_{\mathrm{ALG}}\right] = v^T \sum_\ell \lambda_\ell x^\ell = \alpha v^T x' = \frac{\alpha}{\gamma} v^T x$. Finally, by Lemma 3.1, we have $\mathbf{E}\left[v^T x\right] \geq \frac{1}{64} v^T x^*$, which proves the theorem. □

We replace the optimal LP solution by the suboptimal solution x', because it fulfills stronger monotonicity conditions that make the mechanism truthful in first-order stochastic dominance.

THEOREM 3.3. *Mechanism 2 is truthful in first-order stochastic dominance.*

PROOF. To show truthfulness in first-order stochastic dominance, we consider the influence that a single bidder has on its assigned entity. Due to scaling and decomposition, there are only two possible cases for bidder i. With probability $\frac{\alpha}{\gamma}$, she gets her selection S_i made in Step 1 of the overselling algorithm and pays $\sum_{e \in [m]} a_{e,(i,S_i)} p^i_e$. With probability $1 - \frac{\alpha}{\gamma}$, she receives no entity and valuation, payments and utility are 0. Note that S_i is the entity maximizing $v_{i,S} - \sum_{e \in [m]} a_{e,(i,S)} p^i_e$ over all entities $S \in \mathcal{S}_i$, and α and γ are independent of any bids. Thus, $\mathbf{Pr}\left[v_i(f(v)) - p_i(v) \geq x\right]$ is 1 for $x \leq 0$, it is $\frac{\alpha}{\gamma}$ for $0 < x \leq v_{i,S_i} - \sum_{e \in [m]} a_{e,(i,S_i)} p^i_e$ and 0 otherwise.

Suppose i lies v'_i or, equivalently, demands a different entity S'_i. The prices p^i_e remain unaffected, but the received entity might be inferior. As the demanded entity is the only one that i is able to obtain,

$$\mathbf{Pr}\left[v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i}) \geq x\right] = 0 < \frac{\alpha}{\gamma} = \mathbf{Pr}\left[v_i(f(v)) - p_i(v) \geq x\right] \quad ,$$

for $v_{i,S'_i} - \sum_{e \in [m]} a_{e,(i,S'_i)} p^i_e < x \leq v_{i,S_i} - \sum_{e \in [m]} a_{e,(i,S_i)} p^i_e$ and

$$\mathbf{Pr}\left[v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i}) \geq x\right] = \mathbf{Pr}\left[v_i(f(v)) - p_i(v) \geq x\right]$$

otherwise. Thus, truthfulness is first-order stochastically dominant. □

In addition to truthfulness, observe that the use of demand oracles ensures individual rationality and no positive transfers.

Analysis of Algorithm 1. It remains to show part (3) of Lemma 3.1: For the LP solution x returned by Algorithm 1 we have $\mathbf{E}\left[v^T x\right] \geq \frac{1}{64} v^T x^*$. For a set $B \subseteq [n]$ of bidders, let

$x^*(B)$ be the optimum solution of the LP where only bidders in B have non-zero entries. Furthermore, let i^{1st} and i^{2nd} be the bidders with the highest and the second highest maximum bid, respectively. In the following, we condition on the event that $i^{\text{1st}} \notin \text{STAT}$ but $i^{\text{2nd}} \in \text{STAT}$. As $v^T x^*([n] \setminus \{i^{\text{2nd}}\}) \geq \frac{1}{2} v^T x^*$, we have

$$\mathbf{E}\left[v^T x^*([n] \setminus \text{STAT}) \;\middle|\; i^{\text{1st}} \notin \text{STAT}, i^{\text{2nd}} \in \text{STAT}\right] \geq \frac{1}{2} v^T x^*([n] \setminus \{i^{\text{2nd}}\}) \geq \frac{1}{4} v^T x^* \ .$$

Furthermore, we let p_e^* be the final price reached for constraint $e \in [m]$ after running the algorithm, i.e., $p_e^* = p_e^{n+1}$.

LEMMA 3.4. *If $i^{2nd} \in \text{STAT}$, $p_e^* < \infty$ for all $e \in [m]$.*

PROOF. Consider any constraint $e \in [m]$. We are done if there is no $t \in [n]$ such that $\sum_{i \leq t} \sum_{S \in \mathcal{S}_i} a_{e,(i,S)} x_{i,S} \geq \gamma - 4$. Hence, let t be the smallest such number. This definition ensures $\sum_{i < t} \sum_{S \in \mathcal{S}_i} a_{e,(i,S)} x_{i,S} < \gamma - 4$ and therefore $\sum_{i \leq t} \sum_{S \in \mathcal{S}_i} a_{e,(i,S)} x_{i,S} < \gamma - 3$. Furthermore, for each $i > t$, we have $p_e^i \geq p_0 \cdot 2^{\gamma - 3} = p_0 \cdot 2^{\log(4mb)+1} \geq bL$. This means, to get outcome $S \in \mathcal{S}_i$, bidder $i > t$ has to bid at least $a_{e,(i,S)} bL$. As we assume $i^{\text{2nd}} \in \text{STAT}$, there is only bidder i^{1st} whose maximum bid exceeds L. Only this bidder can be allocated an entity S such that $a_{e,(i,S)} \geq \frac{1}{n}$. Therefore, we have $\sum_{i > t} \sum_{e \in [m]} a_{e,(i,S)} x_{i,S} \leq 2$. In combination, this means $\sum_{i,S} a_{e,(i,S)} x_{i,S} < \gamma - 1$ and therefore $p_e^* < \infty$. □

LEMMA 3.5. *If $i^{2nd} \in \text{STAT}$, then $v^T x \geq \sum_{e \in [m]} p_e^* - m p_0$.*

To prove the lemma, we use the following extension of the geometric-sum formula.

CLAIM 3.6. *Let $r > 1$. For all $a_1, a_2, \ldots, a_n \in [0, 1]$, we have*

$$\sum_{i \in [n]} a_i r^{s_{i-1}} \leq \frac{r^{s_n} - 1}{r - 1} \ , \text{ where } s_i = \sum_{j=1}^{i} a_j \ .$$

PROOF. We show this claim by induction on n. For $n = 0$, we have $0 = 0$ and therefore the claim holds. For the induction step, we first observe that $1 + a_{n+1}(r-1) \geq r^{a_{n+1}}$. This is due to the fact that the function $f\colon [0,1] \to \mathbb{R}$, $f(x) = r^x$ is convex. Therefore, we have $r^x = f(x) \leq x \cdot f(1) + (1-x) f(0) = xr + (1-x) = 1 + x(r-1)$. Now, we have

$$\sum_{i \in [n+1]} a_i r^{s_{i-1}} = \sum_{i \in [n]} a_i r^{s_{i-1}} + a_{n+1} r^{s_n} \ .$$

By induction hypothesis, the first part is upper-bounded by $\frac{r^{s_n}-1}{r-1}$, giving us in total an upper bound of

$$\begin{aligned}
\frac{r^{s_n} - 1}{r - 1} + a_{n+1} r^{s_n} &= \frac{r^{s_n} + (r-1) a_{n+1} r^{s_n} - 1}{r - 1} \\
&= \frac{r^{s_n} \left(1 + (r-1) a_{n+1}\right) - 1}{r - 1} \\
&\geq \frac{r^{s_n} r^{a_{n+1}} - 1}{r - 1} \\
&= \frac{r^{s_{n+1}} - 1}{r - 1} \ .
\end{aligned}$$

□

This claim is used in the proof of the lemma as follows.

PROOF OF LEMMA 3.5. Let $\ell_e^* = \sum_{i,S} a_{e,(i,S)} x_{i,S}$. As we assume $i^{\text{2nd}} \in \text{STAT}$, we have $p_e^* < \infty$ for all $e \in [m]$ and therefore $p_e^* = 2^{\ell_e^*} p_0$. For each bidder $i \in [n]$, we have $v_i(S_i) \geq \sum_{e \in [m]} a_{e,(i,S_i)} p_e^i$. and thus

$$v^T x \geq \sum_{i \in [n], S \in \mathcal{S}_i} \sum_{e \in [m]} a_{e,(i,S)} x_{i,S} p_e^i \ .$$

As we have $p_e^{i+1} = 2^{a_{e,(i,S)} x_{i,S}} \cdot p_e^i$, Claim 3.6 yields a lower bound on this sum of

$$p_0 \sum_{e \in [m]} \frac{2^{\ell_e^*} - 1}{2 - 1} = \sum_{e \in [m]} p_e^* - m p_0 \ .$$

□

LEMMA 3.7. *It holds* $v^T x \geq v^T x^*([n] \setminus \text{STAT}) - \sum_{e \in [m]} p_e^*$.

PROOF. For all $i \in [n]$, S_i is set to option S maximizing $v_{i,S} - \sum_{e \in [m]} a_{e,(i,S)} p_e^i$. That is, we have

$$v_{i,S_i} - \sum_{e \in [m]} a_{e,(i,S_i)} p_e^i \geq v_{i,S} - \sum_{e \in [m]} a_{e,(i,S)} p_e^i$$

for all S. As $a_{e,(i,S_i)} p_e^i \geq 0$, $p_e^i \leq p_e^*$, and $a_{e,(i,S_i)} \leq 1$, this implies

$$v_{i,S_i} \geq v_{i,S} - \sum_{e \in [m]} p_e^*$$

for all S. Furthermore, $x^*([n] \setminus \text{STAT})$ is an LP solution. Therefore, we have $\sum_S x_{i,S}^*([n] \setminus \text{STAT}) \leq 1$. In combination, we get

$$\begin{aligned} v_{i,S_i} &\geq \sum_S x_{i,S}^*([n] \setminus \text{STAT}) v_{i,S_i} \\ &\geq \sum_S x_{i,S}^*([n] \setminus \text{STAT}) \left(v_{i,S} - \sum_{e \in [m]} p_e^* \right) \\ &\geq v^T x^*([n] \setminus \text{STAT}) - \sum_{e \in [m]} p_e^* \ . \end{aligned}$$

□

When adding the inequalities shown in the previous two lemmas, we obtain $2 v^T x \geq v^T x^*([n] \setminus \text{STAT}) - m p_0$ if $i^{\text{2nd}} \in \text{STAT}$. If furthermore $i^{\text{1st}} \notin \text{STAT}$, we have $L \leq v^T x^*([n] \setminus \text{STAT})$ and therefore $m p_0 = \frac{L}{4} \leq \frac{v^T x^*([n] \setminus \text{STAT})}{4}$. That is, we have

$$\mathbf{E}\left[v^T x \mid i^{\text{1st}} \notin \text{STAT}, i^{\text{2nd}} \in \text{STAT} \right] \geq \mathbf{E}\left[\frac{1}{4} v^T x^*([n] \setminus \text{STAT}) \,\middle|\, i^{\text{1st}} \notin \text{STAT}, i^{\text{2nd}} \in \text{STAT} \right] \ .$$

In total, we get

$$\begin{aligned} &\mathbf{E}\left[v^T x \right] \\ &\geq \mathbf{Pr}\left[i^{\text{1st}} \notin \text{STAT}, i^{\text{2nd}} \in \text{STAT} \right] \mathbf{E}\left[v^T x \mid i^{\text{1st}} \notin \text{STAT}, i^{\text{2nd}} \in \text{STAT} \right] \\ &\geq \mathbf{Pr}\left[i^{\text{1st}} \notin \text{STAT}, i^{\text{2nd}} \in \text{STAT} \right] \frac{1}{4} \mathbf{E}\left[v^T x^*([n] \setminus \text{STAT}) \mid i^{\text{1st}} \notin \text{STAT}, i^{\text{2nd}} \in \text{STAT} \right] \end{aligned}$$

$$\geq \frac{1}{64} v^T x^* \ .$$

4. APPLICATIONS

Using the techniques we present in this paper many existing mechanisms can be strengthened. In this section, we give some example applications considering problems for which so far only truthful-in-expectation mechanisms were known. There is a large number of mechanisms that use the framework by Lavi and Swamy, for instance, mechanisms for geometric intersection problems [Christodoulou et al. 2010] or secondary spectrum auction [Hoefer et al. 2014; Hoefer and Kesselheim 2012; Gopinathan et al. 2011; Zhu et al. 2012]. We here explain the necessary adjustments for two classic problems concerning routing in graphs. In either setting, we are given a graph $G = (V, E)$ with m edges and edge capacities $u_e \geq 1$. The bidders wish to have source and sink nodes connected. In the *edge-disjoint paths* problems (EDP), source node s_i needs to be connected to t_i by a path, while at most u_e paths may use the edge e at the same time. The *all-or-noting multi-commodity flow* problem (ANF) is a relaxation, where instead of a path one only needs a flow of value one to connect s_i to t_i. Again, the maximum flow over edge e has to be bounded by u_e.

For a number of variants of EDP and ANF, Lavi and Swamy [2011] present truthful-in-expectation mechanisms that are built upon approximation algorithms by Chekuri et al. [2005]. All of these mechanisms are for the "known" case, in which each bidder wants to have one source-sink pair connected, which is public knowledge. So the only private information is the respective valuation w_i for being connected. This makes this problem a single-parameter scenario and our transformation from Section 2 can directly be applied. At this point, it is important that the payments can be computed in polynomial time by binary search. Therefore, we get efficient mechanisms that are truthful in first-order stochastic dominance and have the following welfare approximation guarantees: (i) $O(\log m)$ for EDP on planar graphs when $B = \min_{e \in E} u_e \geq 2$, (ii) $O(\log^2 m)$ for ANF on general graphs, and $O(\log m)$ for planar graphs.

Going beyond the single-parameter case, e.g., if sources and sinks are private information, the mechanisms designed in [Lavi and Swamy 2011] do not necessarily ensure truthfulness. Fortunately, our multi-parameter mechanism can cover this case, too; even when a bidder has multiple options of source and sink nodes with different valuations. To see this, we first realize that the LP (Route-P) in [Lavi and Swamy 2011] fits into the pattern described in Section 3. The matrix A has $|E|$ rows, and there is a polynomial-time demand oracle solving only shortest-path problems. By applying Algorithm 1, we get an LP solution with an important property: For each bidder there is at most one path selected. For this reason, we can perform the rounding the same way as in the single-parameter case. Truthfulness in first-order stochastic dominance is ensured by the fact that any bidder gets her preferred paths connected with probability exactly $\frac{\alpha}{\gamma}$. Without making any further assumptions, we lose a factor of at most $O(\log|E| + \log n)$ in comparison to the social-welfare guarantees described above.

In all of the other application domains mentioned above, we can apply our transformation from Section 2 to the single-parameter settings, where our lottery payments can be computed in polynomial time by binary search. This way, the truthfulness is strengthened to truthfulness in first-order stochastic dominance. The multi-parameter problems fit into the framework described in Section 3. As the involved LPs have only polynomially many constraints in the input size, one loses at most a logarithmic factor in the input size in comparison to the known approximation results for truthfulness in expectation.

5. CRITICAL DISCUSSION OF OUR APPROACH

A key aspect in our approach is that we make no assumptions on the particular risk attitude of the bidders. In our approach, bidders are comparing lottery outcomes with respect to stochastic dominance over wealth $v_i - p_i$. One might equivalently assume that they are expectation maximizers over $u_i(v_i - p_i)$, where $u_i : \mathbb{R} \to \mathbb{R}$ is a monotone function that encodes the attitude towards risk. From this perspective, we ensure that the outcome maximizes expectation over $u_i(v_i - p_i)$ for *every monotone function* u_i – including, in particular, linear ones.

For second-order stochastic dominance, we assume only slightly more information by restricting to all classes of risk-averse bidders – including risk-neutral ones. This implies that the outcome maximizes expectation over $u_i(v_i - p_i)$ for *every monotone and concave function* u_i – including linear ones. In this way, the sets of truthful-in-stochastic-dominance mechanisms obviously must be refinements of the set of truthful-in-expectation mechanisms. This is in contrast to recent approaches in the literature about mechanisms in the Bayesian setting, where it is assumed that mechanisms can be tailored towards particular classes of functions u_i, which then gives more power to the mechanism designer and allows to, e.g., extract much more revenue.

Our approach addresses randomized mechanism design in the prior-free setting where it is common to maximize social welfare given by the sum of valuations $\sum_i v_i$. In this way, we measure the total monetary value generated by the allocation, without payment money that stays within the system. While this has a natural appeal when thinking about bidders comparing lotteries, it is less obvious to take this approach when thinking about bidders optimizing concave utility functions. The attitude towards risk can be seen as a distortion of incentives that happens due to the randomization of the mechanism. An alternative definition of social welfare [Dughmi and Peres 2012; Fu et al. 2013] is to take this distortion into account and apply a utilitarian approach: $\sum_i p_i + \sum_i u_i(v_i - p_i)$. Here the first part is the utility (i.e., the revenue) of the (risk-neutral) principal.

Which of these two notions of social welfare is more reasonable? Our approach to social welfare can be seen as an "ex-post" measure of wealth generated for society. After the mechanism has made all coin flips, no bidder is subject to risk and thus values the outcome by $v_i - p_i$. Hence, the expected total wealth in the system after the mechanism is executed is given exactly by the sum of valuations. Instead, the latter approach taking the risk attitude into account can be seen as an "ex-ante" approach, in which the bidders perception of risk due to random choices of the mechanism is incorporated.

We believe both approaches have their merits. In particular, maximizing ex-post wealth of the society is a desirable goal for a mechanism. It naturally extends the definition of social welfare of deterministic mechanism design. Randomization is only used to cope with computational limitations. In contrast, maximizing the ex-ante utilities inherently requires randomization and knowledge of attitudes towards risk. Eliciting utility functions in a prior-free mechanism would, e.g., require an extension of the bidding space. This might be unnatural, impractical, and give rise to additional issues, which can be avoided in our approach.

6. CONCLUSIONS AND OPEN PROBLEMS

In this paper, we have studied truthfulness concepts based on stochastic dominance for mechanisms with money. These concepts help to incentivize agents that are sensitive to risk. Our main results are a black-box transformation for single-parameter domains and an application of the randomized meta-rounding technique to yield polynomial-time mechanisms for multi-dimensional packing domains. There are a large number of open problems, for instance, the existence of black-box transformations for meaningful domains. More generally, it is important to obtain separation results to the standard concepts truthfulness in expectation and universal truthfulness in order to understand in which way stochastic dominance gives different power to a mechanism designer. Last but not least, the design of good mechanisms approximating social desiderata such as welfare or revenue in combination with truthfulness in stochastic dominance is an interesting field of future work.

REFERENCES

ABDULKADIROGLU, A. AND SÖNMEZ, T. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica 66,* 3, 689–701.

ARCHER, A. AND TARDOS, É. 2001. Truthful mechanisms for one-parameter agents. In *Proc. 42nd Symp. Foundations of Computer Science (FOCS)*. 482–491.

AZIZ, H., BRANDT, F., AND BRILL, M. 2013. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In *Proc. 12th Conf. Autonomous Agents and Multi-Agent Systems (AAMAS)*. 287–294.

BHALGAT, A., CHAKRABORTY, T., AND KHANNA, S. 2012. Mechanism design with risk aversion. In *Proc. 8th Intl. Workshop Internet & Network Economics (WINE)*. 198–211.

BOGOMOLNAIA, A. AND MOULIN, H. 2001. A new solution to the random assignment problem. *J. Econom. Theory 100*, 295–328.

BOGOMOLNAIA, A. AND MOULIN, H. 2004. Random matching unter dichotomous preferences. *Econometrica 72*, 257–279.

CHEKURI, C., KHANNA, S., AND SHEPHERD, B. 2005. Multicommodity flow, well-linked terminals, and routing problems. In *Proc. 37th Symp. Theory of Computing (STOC)*. 183–192.

CHRISTODOULOU, G., ELBASSIONI, K., AND FOUZ, M. 2010. Truthful mechanisms for exhibitions. In *Proc. 6th Intl. Workshop Internet & Network Economics (WINE)*. 170–181.

DOBZINSKI, S. AND DUGHMI, S. 2013. On the power of randomization in algorithmic mechanism design. *SIAM J. Comput. 42,* 6, 2287–2304.

DUGHMI, S. 2012. Randomization and computation in strategic settings. Ph.D. thesis, Stanford University.

DUGHMI, S. AND PERES, Y. 2012. Mechanisms for risk averse agents, without loss. CoRR 1206.2957.

DUGHMI, S., ROUGHGARDEN, T., AND YAN, Q. 2011. From convex optimization to randomized mechanims: Toward optimal combinatorial auctions. In *Proc. 43rd Symp. Theory of Computing (STOC)*. 149–158.

EHLERS, L., PETERS, H., AND STORCKEN, T. 2002. Strategy-proof probabilistic decision schemes for one-dimensional single-peaked preferences. *J. Econom. Theory 105*, 408–434.

ESÖ, P. AND FUTO, G. 1999. Auction design with a risk averse seller. *Econom. Lett. 65,* 1, 71–74.

FU, H., HARTLINE, J., AND HOY, D. 2013. Prior-independent auctions for risk-averse agents. In *Proc. 14th Conf. Electronic Commerce (EC)*. 471–488.

GIBBARD, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica 45,* 3, 665–681.

GOPINATHAN, A., LI, Z., AND WU, C. 2011. Strategyproof auctions for balancing social welfare and fairness in secondary spectrum markets. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*. 3020–3028.

HOEFER, M. AND KESSELHEIM, T. 2012. Secondary spectrum auctions for symmetric and submodular bidders. In *Proc. 13th Conf. Electronic Commerce (EC)*. 657–671.

HOEFER, M., KESSELHEIM, T., AND VÖCKING, B. 2013. Truthfulness and stochastic dominance with monetary transfers. In *Proc. 14th Conf. Electronic Commerce (EC)*. 567–582.

HOEFER, M., KESSELHEIM, T., AND VÖCKING, B. 2014. Approximation algorithms for secondary spectrum auctions. *ACM Trans. Internet Techn. 14,* 2–3, 16.

Hu, A., Matthews, S., and Zou, L. 2010. Risk aversion and optimal reserve prices in first- and second-price auctions. *J. Econom. Theory 145,* 3, 1188–1202.

Katta, A.-K. and Sethuraman, J. 2006. A solution to the random assignment problem on the full preference domain. *J. Econom. Theory 131*, 231–250.

Krysta, P. and Vöcking, B. 2012. Online mechanism design (randomized rounding on the fly). In *Proc. 39th Intl. Coll. Automata, Languages and Programming (ICALP).* 636–647.

Lavi, R. and Swamy, C. 2011. Truthful and near-optimal mechanism design via linear programming. *J. ACM 58,* 6, 25.

Mas-Colell, A., Whinston, M., and Green, J. 1995. *Microeconomic Theory.* Oxford University Press.

Maskin, E. and Riley, J. 1984. Optimal auctions with risk averse buyers. *Econometrica 52,* 6, 1473–1518.

Moulin, H. 1980. On strategy-proofness and single peakedness. *Public Choice 35,* 4, 437–455.

Myerson, R. 1981. Optimal auction design. *Math. Oper. Res. 6*, 58–73.

Nisan, N. 2007. Introduction to mechanism design (for computer scientists). In *Algorithmic Game Theory*, N. Nisan, É. Tardos, T. Roughgarden, and V. Vazirani, Eds. Cambridge University Press, Chapter 9.

Peters, H., Schulteis, T., and Vermeulen, D. 2010. Generalized stochastic dominance and bad outcome aversion. *Social Choice and Welfare 35,* 2, 285–290.

Sundararajan, M. and Yan, Q. 2010. Robust mechanisms for risk-averse sellers. In *Proc. 11th Conf. Electronic Commerce (EC)*. 139–148.

Vöcking, B. 2012. A universally-truthful approximation scheme for multi-unit auctions. In *Proc. 23rd Symp. Discrete Algorithms (SODA)*. 846–855.

Zhou, L. 1990. On a conjecture by Gale about one-sided matching problems. *J. Econom. Theory 52*, 123–135.

Zhu, Y., Li, B., and Li, Z. 2012. Truthful spectrum auction design for secondary networks. In *Proc. 31st IEEE Conf. Computer Communications (INFOCOM)*. 873–881.