Explorations in Cyber International Relations

Massachusetts Institute of Technology Harvard University

Knightian Auctions

Alessandro Chiesa
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of

Silvio Micali
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of

Zeyuan Allen Zhu
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of

January 10, 2012

This material is based on work supported by the U.S. Office of Naval Research, Grant No. N00014-09-1-0597. Any opinions, findings, conclusions or recommendations therein are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research.

Knightian Auctions

Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu

January 10, 2012

KNIGHTIAN AUCTIONS[1]

ALESSANDRO CHIESA[a], SILVIO MICALI[b], AND ZEYUAN ALLEN ZHU[c]

We study single-good auctions in a setting where each player knows his own valuation only within a constant multiplicative factor $\delta \in (0, 1)$, and the mechanism designer knows δ. The classical notions of implementation in dominant strategies and implementation in undominated strategies are naturally extended to this setting, but their power is vastly different.

On the negative side, we prove that no dominant-strategy mechanism can guarantee social welfare that is significantly better than that achievable by assigning the good to a random player.

On the positive side, we provide tight upper and lower bounds for the fraction of the maximum social welfare achievable in undominated strategies, whether deterministically or probabilistically.

KEYWORDS: Knightian players, mechanism design, auctions.

1. INTRODUCTION

The goal of this paper is to design mechanisms guaranteeing high social welfare in auctions of a single good whose players are *Knightian*.

1.1. *Knightian Players*

In a traditional single-good auction, each player i is assumed to know his true valuation for the good, θ_i, *exactly*. The assumption, however, may be quite strong. For instance, can i be really sure that his true valuation is exactly \$17,975 rather than —say— \$18,001? If not, then *how can his uncertainty be modeled?*

A classical answer is to assume that i knows the single probability distribution D_i from which his true valuation is drawn. More generally, Knight

[1]Supported in part by ONR Grant No. NOOO14-09-0597.

[a] MIT CSAIL; alexch@csail.mit.edu; http://people.csail.mit.edu/alexch/

[b] MIT CSAIL; silvio@csail.mit.edu; http://people.csail.mit.edu/silvio/

[c] MIT CSAIL; zeyuan@csail.mit.edu; http://zeyuan.0sites.org/

(1921), and later on Bewley (2002), suggested to assume that i knows only a *set* of distributions, one of which is guaranteed to be D_i.

In an auction, however, from a strategic perspective a "Knightian player" could collapse each candidate distribution in his set to its expected value. Accordingly, without loss of generality, in a *Knightian auction* each player i only knows a set of integers, K_i, guaranteed to contain his true valuation θ_i. Therefore, K_i is "the set of all possible candidates for θ_i in i's mind", and will be referred to as i's *candidate-valuation set*.

1.2. *Our Knightian Focus*

Knightian players have received much attention in *decision theory* or in mechanisms with a *single* player. We are instead interested in studying the competition of *multiple* Knightian players in *full-fledged* mechanisms. Transforming rich (i.e., exact or Bayesian) knowledge into optimal mechanisms is important. But equally important is to understand whether there are good mechanisms when the players only have set-theoretic knowledge about themselves.

Specifically, we focus on Knightian auctions of a single good, adopting for simplicity sake a finite perspective. Namely,

- all valuations will be integers between 0 and a *valuation bound* B, and
- all mechanisms specify finitely many pure strategies for each player.

(Our results can however be extended to infinite settings as well.)

1.3. *Knightian Mechanism Design*

Intuitively, a mechanism cannot perform well in a Knightian setting where the candidate-valuation sets K_i are too "spread out", but it might perform well when they are sufficiently "clustered". Accordingly, we believe that performance should be measured as a function of the "inaccuracy" of the players' knowledge.

Measuring inaccuracy. For a candidate-valuation set K_i of a player i, we set

$$\delta_i \stackrel{\text{def}}{=} \frac{\max K_i - \min K_i}{\max K_i + \min K_i} \ .$$

Then, it is immediately seen that $\delta_i \in [0,1]$ and that, because $\theta_i \in K_i$, "player i knows θ_i within a multiplicative factor of δ_i".

We refer to δ_i as *i's individual inaccuracy* (about his internal knowledge).

We define the *(global) inaccuracy* of a Knightian setting to be $\delta \stackrel{\text{def}}{=} \max_i \delta_i$.

A setting is *Knightian* if $\delta > 0$ and *traditional* if $\delta = 0$.

(If $c \in \mathbb{R}$ and $\alpha \in [0,1]$, then we may call $[c - \alpha c, c + \alpha c]$ a *α-interval* with *center* c. Any set contained in a α-interval will be called *α-approximate*.)

Designer knowledge. To study auction mechanisms in a Knightian setting we must specify what information is available to the designer.

In a Bayesian setting, where the true valuation profile is drawn from a common prior distribution D, it is traditionally assumed that D itself is known to the designer (and the players). In a Knightian setting, each player i knows a set of distributions $\{D_{i,1}, \ldots, D_{i,k}\}$, which, as already observed, is equivalent to knowing the set of their expected values, K_i. So: how much of this information should the designer be allowed to know?

An extreme assumption is that he knows every K_i. A much weaker assumption is that he knows all individual inaccuracies, but not any candidate-valuation set K_i. A yet weaker assumption is that he knows only the maximum individual inaccuracy. This is the assumption we choose to work under: namely, when designing an auction mechanism for Knightian setting,

the designer only knows the global inaccuracy parameter δ.

Performance. Being in a tough set-theoretic setting, we adopt a worst-case analysis for evaluating the social welfare performance of an auction mechanism M. Figurately speaking, we envisage the following process. First,

an auction mechanism M is announced for selling a given good to n Knightian players, where each player i privately knows his own candidate-valuation set K_i. Then —aware of $K \stackrel{\text{def}}{=} K_1 \times \cdots \times K_n$ and M, and intending to fool M— the devil secretly chooses a true-valuation profile $\theta \in K$. After that, each player i chooses a strategy σ_i. (A player i may learn θ_i only after the auction is over, or never, but during the auction he acts based only on K_i.) Finally, M is played with strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ so as to produce a winner $w = w(M, \sigma)$. Note that w is in general a random variable, since every σ_i may be mixed and M may be probabilistic.

The *maximum social welfare* is of course $\max_i \theta_i$; and the *realized social welfare* of M on σ is $\mathbb{E}[\theta_{w(M,\sigma)}]$. Informally speaking, the *social welfare performance* of M relative to K and σ is taken to be $\min_{\theta \in K} \frac{\mathbb{E}[\theta_{w(M,\sigma)}]}{\max_i \theta_i}$. (Formally, of course, we must specify the solution concept "behind σ".)

Objectives. In designing mechanisms in a Knightian setting we study and try to maximize their performance as a function of the global inaccuracy δ. In essence, δ is our chosen Trojan horse for bringing meaningful mechanism design in the Knightian setting. Without paying attention to the quality of the players' knowledge about themselves, one might design elementary mechanisms, but not "good" ones.

Q&A.

- *Multiplicative or additive accuracy?* A greater level of generality is achieved by considering two distinct global inaccuracy parameters: a multiplicative one, δ^*, and an additive one, δ^+, leading to the following modified constraint: for all player i there exists $x_i \in \mathbb{R}$ such that

$$K_i \subseteq \left[(1-\delta^*)x_i - \delta^+, (1+\delta^*)x_i + \delta^+\right] \cap \{0, \ldots, B\}.$$

  All of our theorems hold for this more general condition. For simplicity, however, we consider only one kind of global inaccuracy parameter, and we find the multiplicative one more meaningful.

- *Can real δ's be really large?* Absolutely. The players' candidate-valuation sets may indeed be "very approximate". Consider a firm participating to an auction for an exclusive license to manufacture solar panels in the US for a period of 25 years. Even if the demand were precisely known in advance, and the only uncertainty came from the firm's ability to lower its costs of production via some breakthrough research, a firm's individual inaccuracy about its own true valuation for the license could easily exceed 0.5.

1.4. *Solution Concepts*

The analysis of every mechanism requires an underlying solution concept. As Knightian settings are settings of *incomplete information* (i.e., settings whose players do not know exactly the true valuations of their opponents), two solution concepts naturally apply: implementation in dominant strategies and implementation in undominated strategies. Of course, both solution concepts need to be properly extended to our setting, but this is naturally done (and in fact done in a way consistent with all prior works).

In essence, a pure strategy s_i of a player i is (very weakly) *Knightian-dominant* if it provides i with a utility at least as large as that of any other strategy t_i of i, no matter what strategies his opponents may choose, and no matter what candidate in K_i may be i's true valuation. A pure strategy s_i of i is *Knightian-undominated* if i does not have any other strategy t_i that

(1) gives i utility at least as great as s_i no matter what strategy subprofiles his opponents may use, and no matter what member of K_i may be i's true valuation, and

(2) gives i utility strictly greater than s_i for at least some strategy subprofile of his opponents and some member of K_i.

(The set of such undominated strategies under a mechanism M is denoted by $\mathsf{UDed}_i^M(K_i)$, or simply by $\mathsf{UDed}_i(K_i)$ when M is clear from context.)

1.5. *Informal Discussion of Our Results*

How much social welfare can we guarantee in auctions? In traditional ones the answer is trivial: 100% in (very-weakly) dominant strategies, via the second-price mechanism. Things are quite different in Knightian auctions.

1.5.1. *Dominant-Strategy Mechanisms Are Meaningful but Inadequate*

Although a Knightian player "does not have a best valuation to bid", very-weakly-dominant strategies continue to be meaningfully defined in a Knightian auction. In a traditional auction the revelation principle (see Myerson (1981)) guarantees that, as far as very-weakly-dominant strategies are concerned, it suffices to consider mechanisms that restrict a player's strategies to (reporting) single valuations. It is easy to see, however, that a natural exentension of the revelation principle continues to apply in Knightian auctions. Specifically, if a very-weakly-dominant strategy mechanism M with a given social welfare performance guarantee exists, then there also exists a *Knightian-direct* mechanism M', with the same performance, where, for every player i, (1) his pure strategy set consists of reporting *sets* of valuations, and (2) truthfully reporting his own candidate-valuation set K_i is very-weakly dominant.

In principle, therefore, there may be a dominant-strategy mechanism that obtains all true candidate-valuation sets, $K_1, \ldots, K_n$, and guarantees a high social welfare performance. Of course, given the inaccuracy of the players' knowledge of their own true valuations, one should expect some degradation of performance relative to the exact-valuation setting. However, one might conjecture that, in a Knightian auction with global inaccuracy δ, a dominant-strategy mechanism might be able to guarantee some δ-dependent fraction —such as $(1-\delta)$, $(1-3\delta)$, or $(1-\delta)^2$— of the maximum social welfare. We prove, however, that even such modest hopes are overly optimistic.

THEOREM 1 (informal) *For all $n \geq 1$, $\delta \in (0,1)$, and $B > \frac{3-\delta}{2\delta}$, no (possibly probabilistic) very-weakly-dominant-strategy-truthful mechanism $M_{n,\delta,B}$ can guarantee a fraction of the maximum social welfare greater than*

$$\frac{1}{n} + \frac{\lfloor \frac{3-\delta}{2\delta} \rfloor + 1}{B}$$

in any Knightian auction with n players, valuation bound B, and inaccuracy parameter δ.

As a *relative* measure of the quality of the players' self knowledge, δ should be independent of the magnitude of the players' valuations. But to ensure an upper bound on the players' valuations, B should be large. Accordingly, the above result essentially implies that any very-weakly-dominant-strategy mechanism can only guarantee a fraction $\approx \frac{1}{n}$ of the maximum social welfare. However, such a fraction can be trivially achieved by the "naive" very-weakly-dominant-strategy mechanism that, dispensing with all bids, assigns the good to a random player! Thus, Theorem 1 essentially says that no dominant-strategy mechanism can be smart: "the optimal one can only be as good as good as the stupid naive one". In other words,

> *dominant strategies are* **intrinsically** *linked to each player having* **exact** *knowledge of either (1) his own valuation, or (2) the* **unique** *possible distribution from which his own valuation has been drawn.*

By showing the limitations of dominant strategies in Knightian auctions, Theorem 1 opens the door to alternative solution concepts: in particular, to implementation in undominated strategies. We actually believe that the Knightian setting will provide a new and vital role for this natural and non-Bayesian implementation notion.

1.5.2. *The Power of Deterministic Undominated-Strategy Mechanisms*

We tightly characterize the power of implementation in undominated strategies via deterministic mechanisms in Knightian auctions. First of all,

without much difficulty, we show that the second-price mechanism (although no longer dominant-strategy) guarantees a relatively good fraction of the maximum social welfare in undominated strategies, despite the fact that it does not leverage any information about δ. Second, more importantly and perhaps more surprisingly, we prove that no deterministic undominated-strategy mechanism can do better, even with full knowledge of δ.

The (good) performance of the second-price mechanism.

THEOREM 2 (informal) *In any Knightian auction with n players, valuation bound B, and inaccuracy parameter δ, the second-price mechanism guarantees a fraction of the maximum social welfare that is*

$$\approx \left(\tfrac{1-\delta}{1+\delta}\right)^2 \quad .^1$$

The course intuition behind Theorem 2 is clear:

> *"It is obvious that each player i should only consider bidding a value v_i inside his own candidate-valuation set K_i. It is further obvious that the worst possible gap between the maximum and the actual social welfare is achieved in the following case. Let w be the winner in the second-price mechanism, and let h, $h \neq w$, be the player with the largest* ***candidate*** *valuation. Player w bids $v_w = \max K_w$, and player h bids $v_h = \min K_h$ (and v_w only slightly exceeds v_h). In this case it is obvious that the second-price mechanism guarantees at most a fraction $\approx \left(\tfrac{1-\delta}{1+\delta}\right)^2$ of the maximum social welfare. Q.E.D."*

Of course, things are a bit more complex. In particular, the fact that a player i should only consider bids in K_i (actually more precisely between $\min K_i - 1$ and $\max K_i + 1$) requires a proof.

[1] We note that when breaking ties at random, the performance of the second-price mechanism is only marginally better: namely, it guarantees a fraction of the maximum social welfare *exactly equal* to $(\frac{1-\delta}{1+\delta})^2$.

The optimality of the second-price mechanism.

THEOREM 3 (informal) *For all $n \geq 2$, $\delta \in (0,1)$, and $B \geq \frac{5}{\delta}$, no deterministic undominated-strategy mechanism $M_{n,B,\delta}$ can guarantee a fraction of the maximum social welfare greater than*

$$\left(\frac{1-\delta}{1+\delta}\right)^2 + \frac{4}{B}$$

in any Knightian auction with n players, valuation bound B, and inaccuracy parameter δ.

Theorem 3 is harder to prove, as is to be expected from an impossibility result. Indeed, its statement applies to all undominated-strategy mechanisms, so that the revelation principle is no longer relevant. Thus, to prove Theorem 3, rather than analyzing a single mechanism (the "direct truthful" one), in principle we should consider *all* possible mechanisms. Considering only those where a player's strategies consist of valuations, or even sets of valuations, is not sufficient. We would have to consider mechanisms with arbitrary strategy sets. Establishing Theorem 3 thus requires new techniques, informally discussed in Section 1.6, and formally provided in Section 6.

1.5.3. *The Greater Power of Probabilistic Undominated-Strategy Mechanisms*

The second-price mechanism "ignores the global inaccuracy parameter". It simply guarantees a fraction $\approx (\frac{1-\delta}{1+\delta})^2$ of the maximum social welfare in any Knightian auction, no matter what the value of δ happens to be. It is thus legitimate to ask whether knowing δ (or a close upper-bound to it) enables one to design mechanisms with better efficiency guarantees. We prove that this is indeed the case: we explicitly construct a probabilistic mechanism that, by properly leveraging δ, outperforms the second-price mechanism, and then we prove that our mechanism is essentially optimal.

THEOREM 4 (informal) *For all $n \geq 2$, $\delta \in (0,1)$, and B, there exists a mechanism $M_{\mathsf{opt}}^{(\delta)}$ that guarantees a fraction of the maximum social welfare that is at least*

$$\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}$$

in any Knightian auction with n players, valuation bound B, and inaccuracy parameter δ.

Theoretical significance. Theorem 4 highlights a novelty of the Knightian setting: namely, probabilism enhances the power of implementation in undominated strategies even for guaranteeing social welfare. By contrast, probabilism offers no such advantage in the exact-valuation world, since the deterministic second-price mechanism already guarantees maximum social welfare. We conjecture that, in Knightian settings, probabilistic mechanisms will enjoy a provably better performance in other applications as well.

Practicality. The proof of Theorem 4 is the technically hardest one in this paper. Nonetheless, we would like to emphasize that $M_{\mathsf{opt}}^{(\delta)}$ is very *practically* played, as it requires almost no computation from the players, and a very small amount of computation from the mechanism. In addition, its performance is *practically* preferable to that of the second-price mechanism. For instance, when $\delta = 0.5$, $M_{\mathsf{opt}}^{(\delta)}$ guarantees a social welfare that is at least **five** times higher than that of the second-price mechanism when there are 2 players, and at least **three** times higher when there are 4 players. (For a full comparison chart, see Appendix A.)

If $M_{\mathsf{opt}}^{(\delta)}$ proves the power of probabilistic mechanisms, our next theorem upperbounds this power by proving that the social-welfare performance of $M_{\mathsf{opt}}^{(\delta)}$ is essentially optimal among all mechanisms, probabilistic or not.

THEOREM 5 (informal) *For all $n \geq 2$, $\delta \in (0,1)$, and $B \geq \frac{5}{\delta}$, no (possibly probabilistic) undominated-strategy mechanism $M_{n,\delta,B}$ can guarantee a fraction of the maximum social welfare greater than*

$$\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} + \frac{4}{B}$$

in any Knightian auction with n players, valuation bound B, and inaccuracy parameter δ.

In sum, our results prove that mechanism design in the Knightian setting is quite possible. Some of the old techniques no longer work, but it is still possible to construct good mechanisms.

1.6. *Two Techniques of Independent Interest*

New ventures require new tools. Let us thus highlight two techniques, crucial to our present endeavor, that we believe will prove useful also to future work in Knightian mechanism design.

The Undominated Intersection Lemma. To prove Theorem 3 and Theorem 5, we establish a basic structural relation between candidate-valuation sets and undominated strategies. The simplest one of course would be $\mathsf{UDed}_i(K_i) = K_i$. This relation, however, is generally false, even when the strategies available to each player consist of individual valuations between 0 and B.[2] A second relation, implied by the previous one, is the following:

$$\forall\ K_i \text{ and } \widetilde{K_i},\ K_i \cap \widetilde{K}_i \neq \emptyset \Rightarrow \mathsf{UDed}_i(K_i) \cap \mathsf{UDed}_i(\widetilde{K}_i) \neq \emptyset.$$

[2]Indeed a mechanism does not need to interpret a bid v_i reported by i as i's true valuation θ_i. For instance, the mechanism could first replace each v_i by $\pi(v_i)$ where π is some fixed permutation over $\{0, 1, \ldots, B\}$ and then run the second-price mechanism as if each player i had bid $\pi(v_i)$. In this case, after $\mathsf{UDed}(K_i)$ has been correctly computed, it will look very different from K_i.

It is not clear, however, whether this second relation always holds.[3] Indeed, an undominated-strategy mechanism may have to specify its strategy sets in quite unforseen ways. Therefore, as soon as K_i and $\widetilde{K}_i$ are even slightly different, their corresponding $\mathsf{UDed}_i(K_i)$ and $\mathsf{UDed}_i(\widetilde{K}_i)$ may in principle be totally unrelated. We prove, however, that the following simple variation of the second relation holds for any possible mechanism. Informally,

> *For any mechanism, probabilistic or not, if K_i and $\widetilde{K}_i$ have at least* **two** *values in common, then there exist two (possibly mixed) "almost payoff-equivalent" strategies σ_i and $\widetilde{\sigma}_i$ respectively having $\mathsf{UDed}_i(K_i)$ and $\mathsf{UDed}_i(\widetilde{K}_i)$ as their support.*

This relation actually suffices for deriving all our impossibility results.

The Distinguishable Monotonicity Lemma. To prove that a given social choice function can be implemented in undominated strategies we are happy to consider mechanisms using a restricted kind of strategies and allocation functions, but we must achieve a delicate balance. On one hand, these restrictions should ensure that the undominated strategies corresponding to a given candidate-valuation set can be characterized in a way that is both *conceptually simple* and *easy to work with*. On the other hand, they should be *sufficient* for proving our Theorem 2 and Theorem 4.

Specifically, we consider mechanisms whose strategies consist of individual valuations (i.e., the pure strategies of each player coincide with $\{0, \ldots, B\}$) and whose allocation functions are restrictions (to $\{0, \ldots, B\}^n$) of integrable functions (over $[0, B]^n$) satisfying a suitable monotonicity property. A simple lemma, the *Distinguishable Monotonicity Lemma*, then guarantees that, for all candidate-valuation set K_i,

$$\mathsf{UDed}_i(K_i) = \{\min K_i, max K_i\}.$$

Although concerned with undominated strategies, when applied to the case

[3] It would actually hold if the total number of coins usable by the players for choosing their mixed strategies were upper-bounded by a fixed constant.

of players knowing their valuations exactly, the Distinguishable Monotonicity Lemma is a strengthening of a classical lemma characterizing (very weakly) dominant-strategy-truthful mechanisms in traditional single-good auctions. Further, the Distinguishable Monotonicity Lemma actually applies to *all* single-parameter domains, not just single-good auctions (the same way that the classical lemma does). We thus believe that this simple lemma will be useful beyond the immediate needs of this paper.

2. PRIOR WORK WITH KNIGHTIAN PLAYERS

As already mentioned, Knightian players have received a lot of attention in decision theory. In particular, Aumann (1962), Dubra et al. (2004), Ok (2002) and Nascimento (2011) investigate decision with incomplete orders of preferences.

The merits of different ways for a Knightian player to "condense" his set of possible values into a meaningfully and deterministically-chosen *single value* have been explored. For example, Danan (2010) studies the average, Gilboa and Schmeidler (1989) the maximum, and Schmeidler (1989) the Choquet expectation.

Other authors have studied mechanisms where a *single* Knightian player is called to accept or reject a given offer; in particular Lopomo et al. (2009) studied the rent-extraction problem in such a setting.

Less relevant to our work, several authors have considered individual Bayesians to model a player's uncertainty: for instance, Sandholm (2000), Porter et al. (2008), and Feige and Tennenholtz (2011). Also, others have studied equilibrium models with unordered preferences: for instance Mas-Colell (1974), Gale and Mas-Colell (1975), Shafer and Sonnenschein (1975), and Fon and Otani (1979). More recently, Rigotti and Shannon (2005) have characterized the set of equilibria in a financial market problem.

3. SINGLE-GOOD KNIGHTIAN AUCTIONS

We separate every auction into two parts, a context and a mechanism.

Knightian Contexts. A Knightian context C has the following components.[4]

- $[n] = \{1, 2, \ldots, n\}$, the set of *players.*
- $\{0, 1, \ldots, B\}$, the set of all *valuations*, where B is the *valuation bound.*
- $\delta \in (0, 1]$, the *inaccuracy* of the context.
- K, the profile of *candidate-valuation sets*, where, for all i, $K_i \subseteq \delta[x_i]$ for some $x_i \in \mathbb{R}$. Here, $\delta[x] \stackrel{\text{def}}{=} [(1-\delta)x, (1+\delta)x] \cap \{0, 1, \ldots, B\}$.
- θ, the profile of *true valuations*, where each $\theta_i \in K_i$.
- $\Omega = \{0, 1, \ldots, n\} \times \mathbb{R}^n$, the set of *outcomes.* If $(a, P) \in \Omega$, then we refer to a as an *allocation* and to P as a profile of *prices.* (If $a = 0$ then the good remains unallocated, else player a wins the good.)
- u, the profile of *utility functions.* Each u_i maps any outcome (a, P) to $\theta_i - P_i$ if $a = i$, and to $-P_i$ otherwise.

Notice that C is fully specified by n, B, δ, K, and θ, that is, $C = (n, B, \delta, K, \theta)$.

KNOWLEDGE. In a context $C = (n, B, \delta, K, \theta)$ each player i only knows K_i and that $\theta_i \in K_i$, and a mechanism designer only knows n, B, and δ.

NOTATION. The set of all contexts with n players, valuation bound B, and inaccuracy δ, is denoted by $\mathscr{C}_{n,B,\delta}$.

The *social welfare* of an outcome (a, P) relative to true-valuation profile θ, $\text{SW}(\theta, (a, P))$, is defined to be θ_a. The *maximum* social welfare of a true-valuation profile θ, $\text{MSW}(\theta)$, is defined to be $\max_{i \in [n]} \theta_i$.

Mechanisms. Our mechanisms for Knightian contexts are finite and ordinary. Indeed a mechanism for $\mathscr{C}_{n,B,\delta}$ is a pair $M = (S, F)$ where

[4]The players' beliefs are not part of our contexts because they cannot affect our strong solution concept.

- $S = S_1 \times \cdots \times S_n$ is the set of all *pure strategy profiles* of M, and
- $F\colon S \to \{0, 1, \ldots, n\} \times \mathbb{R}^n$ is M's *outcome function*.

Set S is always finite and non-empty, and function F may be probabilistic.

NOTATION.

- We denote pure strategies by Latin letters, and possibly mixed strategies by Greek ones.
- If $M = (S, F)$ is a mechanism and $s \in S$, then by $F_i^A(s)$ and $F_i^P(s)$ we respectively denote the probability that the good is assigned to player i and the expected price paid by i under strategy profile s. For mixed strategy profile $\sigma \in \Delta(S)$, we define $F_i^A(\sigma) \stackrel{\text{def}}{=} \mathbb{E}_{s \leftarrow \sigma}\big[F_i^A(s)\big]$ and $F_i^P(\sigma) \stackrel{\text{def}}{=} \mathbb{E}_{s \leftarrow \sigma}\big[F_i^P(s)\big]$, where $s \leftarrow \sigma$ denotes that "s is drawn from the mixed strategy σ".
- We refer to F^A as the *allocation function of* M. More generally, we say that $f\colon S \to [0,1]^n$ is an *allocation function* if for all strategy profile $s \in S$, $\sum_{i \in [n]} f_i(s) \leq 1$.

4. KNIGHTIAN DOMINANCE

In extending the three classical notions of dominance to our approximate valuation setting, the obvious constraint is that when each candidate-valuation set K_i consists of a single element, then all extended notions must collapse to the original ones.

DEFINITION 4.1 *In a mechanism $M = (S, F)$ for $\mathscr{C}_{n,B,\delta}$, let K be a profile of candidate-valuation sets, i a player, σ_i a (possibly mixed) strategy of i, and s_i a pure strategy of i. Then, relatively to K_i, we say that*

- σ_i **very-weakly dominates** s_i, *in symbols* $\sigma_i \overset{\text{vw}}{\underset{i,K_i}{\succ}} s_i$, *if* $\forall\, \theta_i \in K_i\,, \forall\, t_{-i} \in S_{-i} \,:\, \mathbb{E}u_i(\theta_i, F(\sigma_i, t_{-i})) \geq \mathbb{E}u_i(\theta_i, F(s_i, t_{-i}))$.
- σ_i **weakly dominates** s_i, *in symbols* $\sigma_i \overset{\text{w}}{\underset{i,K_i}{\succ}} s_i$, *if* (*a*) $\sigma_i \overset{\text{vw}}{\underset{i,K_i}{\succ}} s_i$ *and* (*b*) $\exists\, \theta_i \in K_i\,, \exists\, t_{-i} \in S_{-i} \,:\, \mathbb{E}u_i(\theta_i, F(\sigma_i, t_{-i})) > \mathbb{E}u_i(\theta_i, F(s_i, t_{-i}))$.

The **(very-weakly-)dominant strategies** *for K_i and K respectively are*

$$\mathsf{Dnt}_i(K_i) \stackrel{\text{def}}{=} \left\{ s_i \in S_i : \forall\, t_i \in S_i\,,\; s_i \underset{i,K_i}{\overset{\text{vw}}{\succ}} t_i \right\} \textit{ and } \mathsf{Dnt}(K) \stackrel{\text{def}}{=} \mathsf{Dnt}_1(K_1)\times\cdots\times\mathsf{Dnt}_n(K_n)\ .$$

The **undominated strategies** *for K_i and K respectively are*

$$\mathsf{UDed}_i(K_i) \stackrel{\text{def}}{=} \left\{ s_i \in S_i : \nexists\, \sigma_i \in \Delta(S_i)\,,\; \sigma_i \underset{i,K_i}{\overset{\text{w}}{\succ}} s_i \right\} \textit{ and } \mathsf{UDed}(K) \stackrel{\text{def}}{=} \mathsf{UDed}_1(K_1)\times\cdots\times\mathsf{UDed}_n(K_n).$$

We use the notation Dnt_i instead of, say, "VWDnt_i" because we have no need to define weakly-dominant or strictly-dominant strategies in a Knightian setting. (Indeed, Theorem 1 shows that even *very*-weakly-dominant strategies cannot guarantee any non-trivial performance.)

We use the notation UDed_i instead of, say, "UWDed_i", because (as in the classical setting) implementation in undominated strategies is defined for weak dominance.[5]

The above extensions of the classical notions are quite straightforward. Only to the extension of weak dominance might require some attention.[6]

Finally, let us note that the following obviously holds.

FACT 4.2 $\mathsf{UDed}_i(K_i) \neq \emptyset$ *for all* K_i.

[5]Also, when considering instead *very-weak* dominance, two "equivalent" strategies may eliminate each other and the set UDed may become empty.

[6] Consider defining "σ_i weakly dominates s_i" using the following alternative quantifications in the additional condition for weak dominance: (1) $\forall\theta_i\forall t_{-i}$, (2) $\exists\theta_i\forall t_{-i}$, and (3) $\forall\theta_i\exists t_{-i}$. Alternatives 1 and 2 do not yield the classical notion of weak dominance when K_i is singleton. Alternative 3 fails to capture the "weakest condition" for which, in absence of special beliefs, a strategy s_i should be discarded in favor of σ_i. Indeed, since we already know that σ_i very-weakly dominates s_i, for player i to discard strategy s_i in favor of σ_i, it should suffice that s_i is strictly worse than σ_i for a *single* true-valuation candidate $\theta_i \in K_i$. That is, we should not insist that s_i be strictly worse than σ_i for all $\theta_i \in K_i$.

5. FORMAL STATEMENT AND PROOF OF THEOREM 1

Theorem 1. *For all $n \geq 1$, $\delta \in (0,1)$, $B > \frac{3-\delta}{2\delta}$, and all (possibly probabilistic) very-weakly-dominant-strategy-truthful mechanisms $M = (S, F)$ for $\mathscr{C}_{n,B,\delta}$, there exists a context $(n, B, \delta, K, \theta) \in \mathscr{C}_{n,B,\delta}$ such that*

$$\mathbb{E}\Big[\mathrm{SW}\big(\theta, F(K)\big)\Big] \leq \left(\frac{1}{n} + \frac{\lfloor \frac{3-\delta}{2\delta} \rfloor + 1}{B}\right) \mathrm{MSW}(\theta) \ .$$

Proof. Fix arbitrarily n, δ, and B such that $B > \frac{3-\delta}{2\delta}$. We start by proving a separate claim. Essentially, as soon as a player reports a δ-interval whose center is sufficiently high, his winning probability remains constant. (Actually the same holds for his price, although we do not care about it.)

Claim 5.1 *For all players i, all integers $x \in (\frac{3-\delta}{2\delta}, B]$, and all subprofiles $\widetilde{K}_{-i}$ of δ-approximate candidate-valuation sets,*

$$F_i^A(\delta[x], \widetilde{K}_{-i}) = F_i^A(\delta[x+1], \widetilde{K}_{-i}).$$

Proof of Claim 5.1: Because K_i may coincide with $\delta[x]$, and because when this is the case reporting $\delta[x]$ very-weakly dominates reporting $\delta[x+1]$, the following inequality must hold: $\forall\, \theta_i \in \delta[x]$,

(5.1)
$$F_i^A(\delta[x], \widetilde{K}_{-i})\cdot\theta_i - F_i^P(\delta[x], \widetilde{K}_{-i}) \geq \ F_i^A(\delta[x+1], \widetilde{K}_{-i})\cdot\theta_i - F_i^P(\delta[x+1], \widetilde{K}_{-i})$$

Because K_i may coincide with $\delta[x+1]$, and because when this is the case reporting $\delta[x+1]$ very-weakly dominates reporting $\delta[x]$, the following inequality also holds: $\forall\, \theta_i' \in \delta[x+1]$,

(5.2)
$$F_i^A(\delta[x+1], \widetilde{K}_{-i})\cdot\theta_i' - F_i^P(\delta[x+1], \widetilde{K}_{-i}) \geq \ F_i^A(\delta[x], \widetilde{K}_{-i})\cdot\theta_i' - F_i^P(\delta[x], \widetilde{K}_{-i}) \ .$$

Thus, setting $\theta_i = x$ in Eq. 5.1 and $\theta'_i = x+1$ in Eq. 5.2, and summing up (the corresponding terms of) the resulting inequalities, the F_i^P price terms and a few other terms cancel out yielding the following inequality:

$$F_i^A(\delta[x+1], \widetilde{K}_{-i}) \geq F_i^A(\delta[x], \widetilde{K}_{-i}) \ . \tag{5.3}$$

Also, setting $\theta_i = \lfloor x(1+\delta) \rfloor$ in Eq. 5.1 and $\theta'_i = \lceil (x+1)(1-\delta) \rceil$ in Eq. 5.2,[7] and summing up the resulting inequalities we obtain the following one:

$$\Big(F_i^A(\delta[x], \widetilde{K}_{-i}) - F_i^A(\delta[x+1], \widetilde{K}_{-i})\Big) \cdot \Big(\lfloor x(1+\delta) \rfloor - \lceil (x+1)(1-\delta) \rceil\Big) \geq 0 \ . \tag{5.4}$$

Now notice that $\lfloor x(1+\delta) \rfloor - \lceil (x+1)(1-\delta) \rceil > 0$, because, by hypothesis, $x > \frac{3-\delta}{2\delta}$. Thus from Eq. 5.4 we deduce

$$F_i^A(\delta[x], \widetilde{K}_{-i}) \geq F_i^A(\delta[x+1], \widetilde{K}_{-i}) \tag{5.5}$$

Together, Eq. 5.3 and Eq. 5.5 imply our claim. *Q.E.D.*

Let us now finish the proof of Theorem 1. Choose the profile of candidate-valuation sets $\widehat{K} \stackrel{\text{def}}{=} (\delta[c], \delta[c], \ldots, \delta[c])$, where $c \stackrel{\text{def}}{=} \lfloor \frac{3-\delta}{2\delta} \rfloor + 1$. By averaging, because the summation of $F_i^A(\widehat{K})$ over $i \in [n]$ cannot be greater than 1, there must exist a player j such that $F_j^A(\widehat{K}) \leq 1/n$. Without loss of generality, let such player be player 1. Then, invoking Claim 5.1 multiple times we have

$$F_1^A(\delta[B], \delta[c], \ldots, \delta[c]) = F_1^A(\delta[B-1], \delta[c], \ldots, \delta[c]) = \cdots = F_1^A(\delta[c], \delta[c], \ldots, \delta[c]) = F_1^A(\widehat{K}) \leq \frac{1}{n} \ .$$

Now suppose that the true candidate-valuation profile of the players is $K \stackrel{\text{def}}{=} (\delta[B], \delta[c], \ldots, \delta[c])$. Then, $\theta = (B, c, \ldots, c) \in K$ and

$$\mathbb{E}\big[\mathrm{SW}\big(\theta, F(K)\big)\big] \leq \frac{1}{n}B + \frac{n-1}{n}c \leq \left(\frac{1}{n} + \frac{c}{B}\right) B = \left(\frac{1}{n} + \frac{c}{B}\right) \cdot \mathrm{MSW}(\theta) \ ,$$

as desired. ■

[7] The hypothesis $x > \frac{3-\delta}{2\delta}$ implies that $x > \frac{1}{2\delta}$, which in turn implies that, under the above choices, $\theta_i \in \delta[x]$ and $\theta'_i \in \delta[x+1]$.

6. THE UNDOMINATED INTERSECTION LEMMA

LEMMA 6.1 (**Undominated Intersection Lemma**) *Let $M = (S, F)$ be a mechanism, i a player, and K_i and $\widetilde{K}_i$ two candidate-valuation sets of i such that $|K_i \cap \widetilde{K_i}| > 1$. Then, for every $\varepsilon > 0$, there exist mixed strategies $\sigma_i \in \Delta(\mathsf{UDed}_i(K_i))$ and $\widetilde{\sigma}_i \in \Delta(\mathsf{UDed}_i(\widetilde{K}_i))$ such that $\forall\, s_{-i} \in S_{-i}$,*

$$\left|F_i^A(\sigma_i, s_{-i}) - F_i^A(\widetilde{\sigma}_i, s_{-i})\right| < \varepsilon \ .$$

(Actually the same holds for F^P, although we do not care about it.)

Proof. Let x_i and y_i be two distinct integers in $K_i \cap \widetilde{K}_i$, and, without loss of generality, let $x_i > y_i$.

Recall that, by Fact 4.2, $\mathsf{UDed}_i(K_i)$ and $\mathsf{UDed}_i(\widetilde{K}_i)$ are both nonempty. If there exists a common (pure) strategy $s_i \in \mathsf{UDed}_i(K_i) \cap \mathsf{UDed}_i(\widetilde{K}_i)$, then setting $\sigma_i = \widetilde{\sigma}_i = s_i$ completes the proof. Therefore, let us assume that $\mathsf{UDed}_i(K_i)$ and $\mathsf{UDed}_i(\widetilde{K}_i)$ are disjoint, and let s_i be a strategy in $\mathsf{UDed}_i(K_i)$ but not in $\mathsf{UDed}_i(\widetilde{K}_i)$. The finiteness of the strategy set S_i implies the existence of a strategy $\widetilde{\sigma}_i \in \Delta(\mathsf{UDed}_i(\widetilde{K}_i))$ such that $\widetilde{\sigma}_i \overset{\mathrm{w}}{\underset{i,\widetilde{K}_i}{\succ}} s_i$.[8] We now prove that

$$\exists\, \tau_i \in \Delta(\mathsf{UDed}_i(K_i)) \text{ such that } \tau_i \overset{\mathrm{w}}{\underset{i,K_i}{\succ}} \widetilde{\sigma}_i.^{9}$$

Let $\widetilde{\sigma}_i = \sum_{j \in X} \alpha^{(j)} \widetilde{s}_i^{(j)}$, where X is a subset of S_i. Invoking again the disjointness of the two undominated strategy sets, we deduce that for each $j \in X$ there exists a strategy $\tau_i^{(j)} \in \Delta(\mathsf{UDed}_i(K_i))$ such that $\tau_i^{(j)} \overset{\mathrm{w}}{\underset{i,K_i}{\succ}} \widetilde{s}_i^{(j)}$. Then it is easily seen that $\tau_i \overset{\text{def}}{=} \sum_{j \in X} \alpha^{(j)} \tau_i^{(j)}$ is a mixed strategy as desired.

For the same reason, we can also find some $\widetilde{\tau}_i \in \Delta(\mathsf{UDed}_i(\widetilde{K}_i))$ such that $\widetilde{\tau}_i \overset{\mathrm{w}}{\underset{i,\widetilde{K}_i}{\succ}} \tau_i$. Continuing in this fashion, "jumping" back and forth between

[8] When S_i is not finite, we need of course to assume that the mechanism is bounded; see Jackson (1992).

[9] Note that, while we have only defined what it means for a *pure* strategy to be dominated by a possibly mixed one, the definition trivially extends to the case of dominated strategies that are *mixed*, as for "$\tau_i \overset{\mathrm{w}}{\underset{i,K_i}{\succ}} \widetilde{\sigma}_i$" in the case at hand.

$\Delta(\mathsf{UDed}_i(K_i))$ and $\Delta(\mathsf{UDed}_i(\widetilde{K}_i))$, we obtain an infinite chain of not necessarily distinct strategies, $\{\sigma_i^{(k)}, \widetilde{\sigma}_i^{(k)}\}_{k\in\mathbb{N}}$ (where $(\sigma_i^{(1)} = s_i,\ \widetilde{\sigma}_i^{(1)} = \sigma_i,$ and) such that

$$\sigma_i^{(1)} \underset{i,K_i}{\overset{\mathrm{w}}{\prec}} \widetilde{\sigma}_i^{(1)} \underset{i,\widetilde{K}_i}{\overset{\mathrm{w}}{\prec}} \sigma_i^{(2)} \underset{i,K_i}{\overset{\mathrm{w}}{\prec}} \widetilde{\sigma}_i^{(2)} \underset{i,\widetilde{K}_i}{\overset{\mathrm{w}}{\prec}} \cdots$$

Since weak dominance implies very-weak dominance, we have that for all $s_{-i} \in S_{-i}$ and all $k \in \mathbb{N}$:

$$\begin{array}{lll} \forall \widetilde{\theta}_i \in \widetilde{K}_i, & F_i^A(\sigma_i^{(k)}, s_{-i})\widetilde{\theta}_i - F_i^P(\sigma_i^{(k)}, s_{-i}) & \leq \quad F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})\widetilde{\theta}_i - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) \\ \forall \theta_i \in K_i, & F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})\theta_i - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) & \leq \quad F_i^A(\sigma_i^{(k+1)}, s_{-i})\theta_i - F_i^P(\sigma_i^{(k+1)}, s_{-i}) \end{array}$$

Thus, for any $z_i \in K_i \cap \widetilde{K}_i$, setting $\theta_i = \widetilde{\theta}_i = z_i$ we see that, for all $s_{-i} \in S_{-i}$ and for $k = 1, 2, \ldots$

$$\begin{array}{rlcl} & F_i^A(\sigma_i^{(k)}, s_{-i})z_i & - & F_i^P(\sigma_i^{(k)}, s_{-i}) \\ \leq & F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})z_i & - & F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) \\ \leq & F_i^A(\sigma_i^{(k+1)}, s_{-i})z_i & - & F_i^P(\sigma_i^{(k+1)}, s_{-i}) \ . \end{array}$$

That is, we have an infinite and non-decreasing sequence that is bounded from above. (Indeed, $z_i \leq B$, F_i^A ranges between 0 and 1, and each price is non-negative.) Thus, since $x_i, y_i \in K_i \cap \widetilde{K}_i$, for any $\varepsilon' > 0$ there exists a (sufficiently large) k such that

$$\left| F_i^A(\sigma_i^{(k)}, s_{-i})x_i - F_i^P(\sigma_i^{(k)}, s_{-i}) - \left(F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) \right) \right| < \varepsilon'$$

and

$$\left| F_i^A(\sigma_i^{(k)}, s_{-i})y_i - F_i^P(\sigma_i^{(k)}, s_{-i}) - \left(F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})y_i - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) \right) \right| < \varepsilon'$$

Now consider the following two linear functions:

$$g(z) \overset{\text{def}}{=} F_i^A(\sigma_i^{(k)}, s_{-i})z - F_i^P(\sigma_i^{(k)}, s_{-i}) \quad \text{and} \quad h(z) \overset{\text{def}}{=} F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})z - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) \ .$$

We have showed that $|g(x_i) - h(x_i)| < \varepsilon'$ and $|g(y_i) - h(y_i)| < \varepsilon'$. We now use the fact that if two linear functions are close at two points, they must

have similar slopes. In particular,

$$\begin{aligned}|\text{slope}(g) - \text{slope}(h)| &= \left|\frac{g(x_i) - g(y_i)}{x_i - y_i} - \frac{h(x_i) - h(y_i)}{x_i - y_i}\right| \\ &\leq \frac{|g(x_i) - h(x_i)| + |g(y_i) - h(y_i)|}{|x_i - y_i|} < \frac{2\varepsilon'}{|x_i - y_i|} \ .\end{aligned}$$

The proof is complete by taking $\varepsilon' = \varepsilon|x_i - y_i|/2$. ■

7. FORMAL STATEMENT AND PROOF OF THEOREM 3

Theorem 3. *For all $n \geq 1$, $\delta \in (0,1)$, $B \geq \frac{5}{\delta}$, and all* deterministic *mechanisms $M = (S, F)$ for $\mathscr{C}_{n,B,\delta}$, there exist a context $(n, B, \delta, K, \theta) \in \mathscr{C}_{n,B,\delta}$ and a strategy profile $s \in \mathsf{UDed}(K)$ such that:*

$$\text{SW}\big(\theta, F(s)\big) \leq \left(\left(\frac{1-\delta}{1+\delta}\right)^2 + \frac{4}{B}\right) \text{MSW}(\theta) \ .$$

Proof. Choose x and y such that $\lfloor(1+\delta)x\rfloor = B$ and $y \stackrel{\text{def}}{=} \lfloor\frac{(1-\delta)x+2}{1+\delta}\rfloor$. Since $B > \frac{5}{\delta}$, $\lceil(1-\delta)y\rceil$ belongs to $\{0, 1, \ldots, B\}$. Furthermore, recalling that $\delta[x] \stackrel{\text{def}}{=} [(1-\delta)x, (1+\delta)x] \cap \{0, 1, \ldots, B\}$, one can verify that $\delta[x]$ and $\delta[y]$ both contain the two integers $\lceil x(1-\delta)\rceil$ and $\lceil x(1-\delta)\rceil + 1$.

Choose ε such that $\frac{1}{n} + \varepsilon < 1$. Then (the Undominated Intersection) Lemma 6.1 guarantees that

(7.1)

$\forall i \in [n]\ \exists \sigma_i \in \Delta(\mathsf{UDed}_i(\delta[x]))$ and $\sigma_i' \in \Delta(\mathsf{UDed}_i(\delta[y]))$ such that $\forall s_{-i} \in S_{-i}$:

$$\left|F_i^A(\sigma_i, s_{-i}) - F_i^A(\sigma_i', s_{-i})\right| < \varepsilon \ .$$

Now consider the allocation distribution $F^A(\sigma_1', \ldots, \sigma_n')$, where the randomness comes from the mixed strategy profile since M is a deterministic mechanism. Since the good will be assigned with a total probability mass of 1, by averaging, there exists a player j such that $F_j^A(\sigma_1', \ldots, \sigma_n') \leq \frac{1}{n}$:

that is, player j wins the good with probability at most $\frac{1}{n}$. Without loss of generality, let $j = 1$. In particular, there exist $s'_{-1} \in \mathsf{UDed}_2(\delta[y]) \times \cdots \times \mathsf{UDed}_n(\delta[y])$, such that $F_1^A(\sigma'_1, s'_{-1}) \leq \frac{1}{n}$. This together with Eq. 7.1 implies that $F_1^A(\sigma_1, s'_{-1}) \leq \frac{1}{n} + \varepsilon < 1$. In turn, this implies that there exists a *pure* strategy $s_1 \in \mathsf{UDed}_1(\delta[x])$ such that, setting $s \stackrel{\text{def}}{=} (s_1, s'_{-1})$, $F_1^A(s) = 0$.

Now we construct the desired δ-approximate candidate-valuation profile K and the true-valuation profile θ as follows:

$$K \stackrel{\text{def}}{=} \big(\delta[x], \delta[y], \ldots, \delta[y]\big) \quad \text{and} \quad \theta \stackrel{\text{def}}{=} \big(\lfloor(1+\delta)x\rfloor, \lceil(1-\delta)y\rceil, \ldots, \lceil(1-\delta)y\rceil\big) \ .$$

Note that $s \in \mathsf{UDed}(K)$, $\theta \in K$, and $\mathrm{MSW}(\theta) = \lfloor(1+\delta)x\rfloor$. Since $F_1^A(s) = 0$,

$$\begin{aligned}
\mathrm{SW}\big(\theta, F(s)\big) &= \lceil(1-\delta)y\rceil \leq (1-\delta)y + 1 \\
&\leq \frac{(1-\delta)^2 x}{1+\delta} + 3 = \frac{(1-\delta)^2}{(1+\delta)^2}(1+\delta)x + 3 \\
&\leq \frac{(1-\delta)^2}{(1+\delta)^2}\lfloor(1+\delta)x\rfloor + 4 \leq \frac{(1-\delta)^2}{(1+\delta)^2}\lfloor(1+\delta)x\rfloor + \frac{4}{B}\lfloor(1+\delta)x\rfloor \\
&\leq \left(\frac{(1-\delta)^2}{(1+\delta)^2} + \frac{4}{B}\right)\lfloor(1+\delta)x\rfloor = \left(\frac{(1-\delta)^2}{(1+\delta)^2} + \frac{4}{B}\right)\mathrm{MSW}(\theta) \ .
\end{aligned}$$

Thus the theorem holds. ■

8. FORMAL STATEMENT AND PROOF OF THEOREM 5

Theorem 5. *For all $n \geq 1$, $\delta \in (0,1)$, $B \geq \frac{5}{\delta}$, and all (deterministic or probabilistic) mechanisms $M = (S, F)$ for $\mathscr{C}_{n,B,\delta}$, there exist a context $(n, B, \delta, K, \theta) \in \mathscr{C}_{n,B,\delta}$ and a strategy profile $s \in \mathsf{UDed}(K)$ such that*

$$\mathbb{E}\Big[\mathrm{SW}\big(\theta, F(s)\big)\Big] \leq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} + \frac{4}{B}\right)\mathrm{MSW}(\theta) \ . \tag{8.1}$$

Proof. (The first part of the proof closely tracks that of Theorem 3 in Appendix 7.[10])

[10]Very informally, the only differences are that the allocation distribution $F^A(\sigma'_1, \ldots, \sigma'_n)$ now depends also on the "coin tosses of the mechanism", and that one can no longer guarantee the existence of a pure strategy s such that $F_1^A(s) = 0$.

Choose x and y such that $\lfloor (1+\delta)x \rfloor = B$ and $y \stackrel{\text{def}}{=} \lfloor \frac{(1-\delta)x+2}{1+\delta} \rfloor$. Then again $\lceil (1-\delta)y \rceil$ belongs to $\{0, 1, \ldots, B\}$ and $\delta[x]$ and $\delta[y]$ both contain the following two integer points: $\lceil x(1-\delta) \rceil$ and $\lceil x(1-\delta) \rceil + 1$.

Since we always have $\lceil (1-\delta)y \rceil < (1-\delta)y+1$, we can choose $\varepsilon \in (0, 1-\frac{1}{n})$ such that

$$\frac{n-1}{n}\lceil (1-\delta)y \rceil + \varepsilon \lceil (1-\delta)x \rceil - \varepsilon \lceil (1-\delta)y \rceil < \frac{n-1}{n}(1-\delta)y + 1 \ .$$

Then (the Undominated Intersection) Lemma 6.1 guarantees that

(8.2)
$\forall i \in [n]$ there exist $\sigma_i \in \Delta(\mathsf{UDed}_i(\delta[x]))$ and $\sigma'_i \in \Delta(\mathsf{UDed}_i(\delta[y]))$ such that $\forall s_{-i} \in S_{-i}$:

$$\left| F_i^A(\sigma_i, s_{-i}) - F_i^A(\sigma'_i, s_{-i}) \right| < \varepsilon \ .$$

Again consider the allocation distribution $F^A(\sigma'_1, \ldots, \sigma'_n)$. By averaging, there exists some player j such that $F_j^A(\sigma'_1, \ldots, \sigma'_n) \leq \frac{1}{n}$. Thus, by our choice of ε and Eq. 8.2, we have that $F_1^A(\sigma_1, \sigma'_{-1}) \leq \frac{1}{n} + \varepsilon$. This implies that there exists a *pure* strategy profile $s = (s_1, s'_{-1})$ that is in the support of (σ_1, σ'_{-1}) —and thus in $\mathsf{UDed}_1(\delta[x]) \times \mathsf{UDed}_2(\delta[y]) \times \cdots \times \mathsf{UDed}_2(\delta[y])$— such that $F_1^A(s_1, s'_{-1}) \leq \frac{1}{n} + \varepsilon$. Now define

$$K \stackrel{\text{def}}{=} \big(\delta[x], \delta[y], \ldots, \delta[y]\big) \quad \text{and} \quad \theta \stackrel{\text{def}}{=} \big(\lfloor (1+\delta)x \rfloor, \lceil (1-\delta)y \rceil, \ldots, \lceil (1-\delta)y \rceil\big) \ .$$

Notice that $s \in \mathsf{UDed}(K)$, $\theta \in K$, and $\mathrm{MSW}(\theta) = \lfloor (1+\delta)x \rfloor$. We now

show that s, K, and θ satisfy the desired Eq. 8.1:

$$\begin{aligned}
\mathbb{E}\Big[\mathrm{SW}\big(\theta, F(s_1, s'_{-1})\big)\Big] &\leq \left(\frac{1}{n}+\varepsilon\right)\cdot\lfloor(1+\delta)x\rfloor + \left(\frac{n-1}{n}-\varepsilon\right)\cdot\lceil(1-\delta)y\rceil \\
&= \frac{1}{n}\cdot\lfloor(1+\delta)x\rfloor + \frac{n-1}{n}\cdot\lceil(1-\delta)y\rceil + \varepsilon\lceil(1-\delta)x\rceil - \varepsilon\lceil(1-\delta)y\rceil \\
&< \frac{1}{n}\cdot\lfloor(1+\delta)x\rfloor + \frac{n-1}{n}\cdot(1-\delta)y + 1 \\
&\leq \frac{1}{n}\cdot\lfloor(1+\delta)x\rfloor + \frac{n-1}{n}\cdot\frac{(1-\delta)^2 x}{1+\delta} + 3 \\
&< \frac{1}{n}\cdot\lfloor(1+\delta)x\rfloor + \frac{n-1}{n}\cdot\frac{(1-\delta)^2}{(1+\delta)^2}\lfloor(1+\delta)x\rfloor + 4 \\
&< \left(\frac{1}{n}+\frac{n-1}{n}\cdot\frac{(1-\delta)^2}{(1+\delta)^2}+\frac{4}{B}\right)\lfloor(1+\delta)x\rfloor \\
&= \left(\frac{1}{n}+\frac{n-1}{n}\cdot\frac{(1-\delta)^2}{(1+\delta)^2}+\frac{4}{B}\right)\mathrm{MSW}(\theta) \\
&= \left(\frac{(1-\delta)^2+\frac{4\delta}{n}}{(1+\delta)^2}+\frac{4}{B}\right)\mathrm{MSW}(\theta)\ .
\end{aligned}$$

■

9. THE DISTINGUISHABLE MONOTONICITY LEMMA

Let us recall a traditional way to define auction mechanisms from suitable allocation functions.

DEFINITION 9.1 *If $f : [0, B]^n \to [0, 1]^n$ is an integrable[11] allocation function, then we denote by M_f the mechanism (S, F) where $S = \{0, 1, \ldots, B\}^n$ and F is so defined: on input bid profile $v \in S$,*

- *with probability $f_i(v)$ the good is assigned to player i, and*
- *if player i wins, he pays $P_i = v_i - \frac{\int_0^{v_i} f_i(z, v_{-i})\,dz}{f_i(v_i, v_{-i})}$ (and all other players pay $P_j = 0$ for $j \neq i$.)*

REMARK 9.2

[11]Specifically, we require that, for each v_{-i}, the function $f_i(z, v_{-i})$ is integrable with respect to z on $[0, B]$.

- M_f is deterministic if and only if $f(\{0,1,\ldots,B\}^n) \subseteq \{0,1\}^n$.
- For all player i and bid profile v, the expected price $F_i^P(v)$ is equal to $v_i \cdot f_i(v_i, v_{-i}) - \int_0^{v_i} f_i(z, v_{-i})\, dz$.
- We stress that M_f continues to have the discrete strategy space $S = \{0,1,\ldots,B\}^n$. The analysis over a continuous domain for f is only a tool for proving the lemma.
- In the exact-valuation world, it is well known that a single-good auction mechanism M is very-weakly dominant-strategy-truthful if and only if $M = M_f$ for some function f that is (integrable and) *monotonic*, that is, such that each f_i is non-decreasing in the bid of player i for any fixed choice of bids of all other players.

We now slightly strengthen the notion of monotonicity.

DEFINITION 9.3 *Let $f : [0,B]^n \to [0,1]^n$ be a allocation function. For $d \in \{1,2\}$, we say that f is d-***distinguishably monotonic** *(d-DM, for short) if f is integrable, monotonic, and satisfying the following "distinguishability" condition:*

$$\forall\, i \in [n]\,,\ \forall v_i, v_i' \in S_i \text{ s.t. } v_i \le v_i' - d,\ \exists\, v_{-i} \in S_{-i} \quad \int_{v_i}^{v_i'} \big(f_i(z, v_{-i}) - f_i(v_i, v_{-i})\big)\, dz > 0\ .$$

If f is d-DM, we say that M_f is d-DM.

Distinguishability is certainly an additional requirement to monotonicity, but actually is a mild one. (Indeed, the second-price mechanism is 2-DM and, if ties are broken at random, even 1-DM.[12]) Yet, in our Knightian setting, this mild additional requirement is quite useful for "controlling" the

[12]For example, the second-price mechanism with lexicographic tie-breaking is the mechanism M_f where f is defined as follows: $\forall i \in [n]$ and $\forall v \in \{0,\ldots,B\}^n$,

$$f_i(v) \stackrel{\text{def}}{=} \begin{cases} 1, & \text{if (a) } v_i > \max v_{-i} \text{ or (b) } v_i = \max v_{-i} \text{ and } i = \min\{j : v_j = v_i\}; \\ 0, & \text{otherwise.} \end{cases}$$

To see that this mechanism is 2-DM, consider two bids v_i and v_i' of player i that are

undominated strategies of a mechanism, and thus for engineering implementations of desirable social choice functions in undominated strategies.

LEMMA 9.4 (**Distinguishable Monotonicity Lemma**) *If f is a d-DM allocation function, then M_f is such that, for any player i and δ-approximate candidate-valuation profile K,*

$$\begin{aligned} &\mathsf{UDed}_i(K_i) \subseteq \{\min K_i, \ldots, \max K_i\} && \text{if } d = 1, \text{ and} \\ &\mathsf{UDed}_i(K_i) \subseteq \{\min K_i - 1, \ldots, \max K_i + 1\} && \text{if } d = 2. \end{aligned}$$

(Above, $\min K_i$ and $\max K_i$ respectively denote the minimum and maximum integers in K_i.)

Proof. For every $i \in [n]$, let $v_i^\perp \stackrel{\text{def}}{=} \min K_i$ and $v_i^\top \stackrel{\text{def}}{=} \max K_i$. Then, to establish our lemma it suffices to prove that, $\forall i \in [n]$ and $\forall d \in \{1, 2\}$, the following four properties hold:

1. *$v_i^\perp$ very-weakly dominates every $v_i \leq v_i^\perp - d$.*
2. *$v_i^\top$ very-weakly dominates every $v_i \geq v_i^\top + d$.*
3. *There is a strategy sub-profile v_{-i} for which $v_i^\perp$ is strictly better than every $v_i \leq v_i^\perp - d$.*
4. *There is a strategy sub-profile v_{-i} for which $v_i^\top$ is strictly better than every $v_i \geq v_i^\top + d$.*

Proof of Property 1. Fix any (pure) strategy sub-profile $v_{-i} \in S_{-i}$ for the other players and any possible true valuation $\theta_i \in K_i$. Letting $v^\perp = (v_i^\perp, v_{-i})$

---

at least a distance of two apart; by choosing a strategy sub-profile for the other players where the highest bid falls between v_i and v_i', we can ensure that the desired integral is positive. A slightly more refined argument shows that the second-price mechanism breaking ties at random is 1-DM.

and $v = (v_i, v_{-i})$, we prove that

$$\begin{aligned}
&\mathbb{E}\Big[u_i\big(\theta_i, F(v^\perp)\big)\Big] - \mathbb{E}\Big[u_i\big(\theta_i, F(v)\big)\Big] \\
&= \big(f_i(v^\perp) - f_i(v)\big) \cdot \theta_i - \big(F_i^P(v^\perp) - F_i^P(v)\big) \\
&= \big(f_i(v^\perp) - f_i(v)\big) \cdot \theta_i - \left(v_i^\perp \cdot f_i(v^\perp) - \int_0^{v_i^\perp} f_i(z, v_{-i})\, dz - v_i \cdot f_i(v) + \int_0^{v_i} f_i(z, v_{-i})\, dz\right) \\
&= \big(f_i(v^\perp) - f_i(v)\big) \cdot (\theta_i - v_i^\perp) + \int_{v_i}^{v_i^\perp} \big(f_i(z, v_{-i}) - f_i(v)\big)\, dz \ .
\end{aligned}$$

Now note that, since $\theta_i \in K_i$, $\theta_i - v_i^\perp = \theta_i - \min K_i \geq 0$; moreover, by the monotonicity of f, whenever $z \geq v_i$, it holds that $f_i(z, v_{-i}) \geq f_i(v)$. We deduce that $\mathbb{E}u_i\big(\theta_i, F(v^\perp)\big) \geq \mathbb{E}u_i\big(\theta_i, F(v)\big)$. We conclude that $v_i^\perp$ very-weakly dominates v_i.

Proof of Property 2. Analogous to that of Property 1 and omitted.

Proof of Property 3. Due to the d-distinguishable monotonicity of M, $v_i \leq v_i^\perp - d$ implies the existence of a strategy sub-profile v_{-i} making $\int_{v_i}^{v_i^\perp} \big(f_i(z, v_{-i}) - f_i(v)\big)\, dz$ strictly positive. For such v_{-i}, therefore, playing $v_i^\perp$ is strictly better than v_i.

Proof of Property 4. Analogous to that of Property 3 and omitted.

Thus the lemma holds. ■

10. FORMAL STATEMENT AND PROOF OF THEOREM 2

Theorem 2. *Let $M_{\mathsf{2P}} = (S_{\mathsf{2P}}, F_{\mathsf{2P}})$ be the second-price mechanism with any deterministic tie-breaking rule. Then, for all contexts $(n, B, \delta, K, \theta)$ and all strategy profiles $v \in \mathsf{UDed}(K)$:*

$$\mathrm{SW}\big(\theta, F_{\mathsf{2P}}(v)\big) \geq \left(\frac{1-\delta}{1+\delta}\right)^2 \mathrm{MSW}(\theta) - 2\,\frac{1-\delta}{1+\delta} \ .$$

Proof. Since K is a δ-approximate candidate-valuation set, for each player i let x_i be such that $K_i \subseteq \delta[x_i]$. Then, in light of (the Distinguishable

Monotonicity) Lemma 9.4 and the previous observation that F^A_{2P} is a 2-DM allocation function, we have that, for each player i:

$$\mathsf{UDed}_i(x) \subseteq \left\{\lceil (1-\delta)x_i \rceil - 1, \ldots, \lfloor (1+\delta)x_i \rfloor + 1\right\} \ . \tag{10.1}$$

Let i^* be the player with the highest true valuation and j^* the player winning the good under the bid profile v, that is, $\theta_{i^*} = \max_i \theta_i$ and $v_{j^*} = \max_j v_j$.

If $i^* = j^*$ then we are done. If $i^* \neq j^*$, we need to show that θ_{j^*} is not much lower than θ_{i^*}.

From Eq. 10.1 we know that $\lceil (1-\delta)x_{i^*} \rceil - 1 \leq v_{i^*}$ and $v_{j^*} \leq \lfloor (1+\delta)x_{j^*} \rfloor + 1$. Because j^* is the winner, we also know that $v_{i^*} \leq v_{j^*}$. Combining these facts and "removing floors and ceilings" we have $(1-\delta)x_{i^*} \leq (1+\delta)x_{j^*} + 2$; equivalently,

$$x_{j^*} \geq \frac{(1-\delta)}{(1+\delta)} x_{i^*} - 2\frac{1}{(1+\delta)} \ .$$

Since we also know that $\theta_{j^*} \geq (1-\delta)x_{j^*}$ and $(1+\delta)x_{i^*} \geq \theta_{i^*}$, we obtain:

$$\begin{aligned} \mathrm{SW}(\theta, F_{\mathsf{2P}}(v)) = \theta_{j^*} \geq (1-\delta)x_{j^*} \geq (1-\delta)\frac{1-\delta}{1+\delta} x_{i^*} - 2\frac{(1-\delta)}{(1+\delta)} \\ \geq (1-\delta)\frac{1-\delta}{1+\delta}\frac{1}{1+\delta}\theta_{i^*} - 2\frac{(1-\delta)}{(1+\delta)} = \frac{(1-\delta)^2}{(1+\delta)^2}\mathrm{MSW}(\theta) - 2\frac{(1-\delta)}{(1+\delta)} \ . \end{aligned}$$

Thus, the theorem holds. ■

Remark 10.1 If $M_{\mathsf{2P}} = (S_{\mathsf{2P}}, F_{\mathsf{2P}})$ were the second-price mechanism breaking ties at random (assigning a positive probability to each tie), then we can use a proof analogous to the one above, with the only difference being that F^A_{2P} is 1-DM (instead of only 2-DM), and invoking the stronger inclusion of (the Distinguishable Monotonicity) Lemma 9.4, to show the following, stronger lower bound:

$$\mathbb{E}\Big[\mathrm{SW}\big(\theta, F_{\mathsf{2P}}(v)\big)\Big] \geq \frac{(1-\delta)^2}{(1+\delta)^2}\mathrm{MSW}(\theta) \ .$$

11. FORMAL STATEMENT AND PROOF OF THEOREM 4

Theorem 4. *$\forall n$, $\forall \delta \in (0,1)$, and $\forall B$, there exists a mechanism $M^{(\delta)}_{\mathsf{opt}} = (S, F)$ such that for every δ-approximate candidate-valuation profile K, every true-valuation profile $\theta \in K$, and every strategy profile $v \in \mathsf{UDed}(K)$:*

$$\mathbb{E}\Big[\mathrm{SW}\big(\theta, F(v)\big)\Big] \geq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}\right) \mathrm{MSW}(\theta) \ .$$

We break the construction of $M^{(\delta)}_{\mathsf{opt}}$ and its analysis into several steps. At the highest level, in order to leverage our Distinguishable Monotonicity Lemma, and thus choose $M^{(\delta)}_{\mathsf{opt}} = M_{f^{(\delta)}}$ for a suitably chosen 1-DM allocation function $f^{(\delta)}$.

11.1. *Our Allocation Function*

Given δ, we find it natural to choose an allocation function $f^{(\delta)}$ that is *symmetric:* that is, $f^{(\delta)}(v') = f^{(\delta)}(v)$ whenever the profile v' consists of a permutation of the bids in v. In other words, "renaming the players should not change the probability of allocating the good to a given player".

Also, when some of the players' bids are much smaller than others, we find it intuitive to interpret the lower bids as being more likely to come from players with lower valuations. Accordingly, our $f^{(\delta)}$ gives positive probability only to the highest bids. However, when the highest bids are close to each other, we find it hard to "infer" which one has been chosen by the player with the highest true valuation: after all, we are in a Knightian model. Therefore our $f^{(\delta)}$ assigns the good to a randomly chosen high-bidding player. A bit more precisely, our $f^{(\delta)}$ deterministically derives from the players' bids a *threshold*, and probabilistically chooses the winning player only among those bids lying above the threshold. To achieve optimality, however, one must be much more careful in allocating probability mass, and some complexities should be expected. Let us now proceed more formally.

Let D_δ be the always positive quantity defined as follows: for all $\delta \in (0,1)$,

$$D_\delta \stackrel{\text{def}}{=} \left(\frac{1+\delta}{1-\delta}\right)^2 - 1 \ .$$

DEFINITION 11.1 *For all $\delta \in (0,1)$, define $f^{(\delta)}\colon [0,B]^n \to [0,1]^n$ as follows: for every $i \in [n]$ and every $z = (z_1, \ldots, z_n) \in [0,B]^n$*

- *if $z_1 \geq z_2 \geq \cdots \geq z_n$, then*

$$f_i^{(\delta)}(z) \stackrel{\text{def}}{=} \begin{cases} \frac{1}{n} \cdot \frac{n+D_\delta}{n^*+D_\delta} \cdot \frac{z_i(n^*+D_\delta) - \sum_{j=1}^{n^*} z_j}{z_i D_\delta}, & \text{if } i \leq n^*, \\ 0, & \text{if } i > n^*; \end{cases} \quad (11.1)$$

  where n^* is the index in $\{1,2,\ldots,n\}$ such that

$$z_1 \geq \cdots \geq z_{n^*} > \frac{\sum_{j=1}^{n^*} z_j}{n^*+D_\delta} \geq z_{n^*+1} \geq \cdots \geq z_n \ . \quad (11.2)$$

- *else, $f_i^{(\delta)}(z) \stackrel{\text{def}}{=} f_{\pi(i)}^{(\delta)}(z_{\pi(1)}, \ldots, z_{\pi(n)})$ where π is any permutation of the players such that $z_{\pi(1)} \geq \cdots \geq z_{\pi(n)}$.*

We refer to $\frac{\sum_{j=1}^{n^*} z_j}{n^*+D_\delta}$ as the **bid threshold***, to players $1, \ldots, n^*$ as the* **candidate winners***, and to the players $n^*+1, \ldots, n$ as the* **losers***.*

11.2. *Our Allocation Function is Well Defined*

LEMMA 11.2 *$f^{(\delta)}$ is an allocation function.*

Proof. Assume, without loss of generality, that $z_1 \geq z_2 \geq \cdots \geq z_n$.

We first prove that n^* exists and is unique, and begin with its existence.

Note that there exists an integer n' in $[n]$ such that

$$\forall\, i > n', \quad z_i \leq \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta} \ . \quad (11.3)$$

Indeed, Eq. 11.3 vacuously holds for $n' = n$. Now letting n'' be the least such integer, the following two facts hold:

$$\forall\, i > n'', \quad z_i \leq \frac{\sum_{j=1}^{n''} z_j}{n''+D_\delta} \quad \text{and} \quad (11.4)$$

$$(11.5)\quad \exists k \geq n'' \text{ such that } z_k > \frac{\sum_{j=1}^{n''-1} z_j}{n''-1+D_\delta} \ .$$

Because z is non-decreasing, the last inequality implies $z_{n''} > \frac{\sum_{j=1}^{n''-1} z_j}{n''-1+D_\delta}$; equivalently, $z_{n''} > \frac{\sum_{j=1}^{n''} z_j}{n''+D_\delta}$. Invoking again the monotonicity of z, we have

$$(11.6)\quad \forall\, i \leq n'', \quad z_i > \frac{\sum_{j=1}^{n''} z_j}{n''+D_\delta} \ .$$

Thus, 11.3 and 11.6 imply that choosing $n^* = n''$ satisfies Eq. 11.2.

Next, we prove that n^* is unique. Suppose by way of contradiction that there exist two integers $n^\perp$ and $n^\top$, $n^\perp < n^\top$, both satisfying Eq. 11.2. Now define

$$S^\perp \stackrel{\text{def}}{=} \sum_{j=1}^{n^\perp} z_j, \quad S^\top \stackrel{\text{def}}{=} \sum_{j=1}^{n^\top} z_j, \quad S^\Delta \stackrel{\text{def}}{=} S^\top - S^\perp, \quad \text{and } n^\Delta \stackrel{\text{def}}{=} n^\top - n^\perp \ .$$

Invoking Eq. 11.2 with $n^\perp$ and $n^\top$, we deduce that for $i \in \{n^\perp+1, \ldots, n^\top\}$,

$$\frac{S^\perp}{n^\perp+D_\delta} \geq z_i > \frac{S^\top}{n^\top+D_\delta} = \frac{S^\perp+S^\Delta}{n^\perp+n^\Delta+D_\delta} \ .$$

Averaging over all z_i such that $i \in \{n^\perp+1, \ldots, n^\top\}$, we get

$$(11.7)\quad \frac{S^\perp}{n^\perp+D_\delta} \geq \frac{S^\Delta}{n^\Delta} > \frac{S^\perp+S^\Delta}{n^\perp+n^\Delta+D_\delta} \ .$$

Let us now show that the second inequality of Eq. 11.7 contradicts the first:

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\perp+S^\Delta}{n^\perp+n^\Delta+D_\delta} \Leftrightarrow (n^\perp+n^\Delta+D_\delta)S^\Delta > n^\Delta(S^\perp+S^\Delta)$$
$$\Leftrightarrow (n^\perp+D_\delta)S^\Delta > n^\Delta S^\perp \Leftrightarrow \frac{S^\Delta}{n^\Delta} > \frac{S^\perp}{(n^\perp+D_\delta)} \ .$$

The contradiction establishes the uniqueness of n^*.

Finally, to prove that $f^{(\delta)}$ is an allocation function we must argue that (a) $f_i^{(\delta)}(z) \geq 0$ for every i and z, and (b) $\sum_i f_i^{(\delta)}(z) \leq 1$ for every z.

Since we are assuming $z_1 \geq z_2 \geq \cdots \geq z_n$, inequality (a) holds because $f_i^{(\delta)}(z) = 0$ for $i > n^*$ by definition, and because Eq. 11.2 tells us that $z_i(n^* + D_\delta) - \sum_{j=1}^{n^*} z_j > 0$ for $i \leq n^*$.

As for inequality (b), since it is easy to see that $\sum_{j=1}^{n^*} \frac{z_j}{z_i} \geq n^*$, we have

$$\begin{aligned}
\sum_{i=1}^{n} f_i^{(\delta)}(z) &= \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} \cdot \sum_{i=1}^{n^*} \frac{z_i(n^* + D_\delta) - \sum_{j=1}^{n^*} z_j}{z_i D_\delta} \\
&= \frac{1}{n} \cdot \frac{n + D_\delta}{(n^* + D_\delta)D_\delta} \cdot \left(n^*(n^* + D_\delta) - \sum_{i=1}^{n^*} \sum_{j=1}^{n^*} \frac{z_j}{z_i} \right) \\
&\leq \frac{1}{n} \cdot \frac{n + D_\delta}{(n^* + D_\delta)D_\delta} \cdot (n^*(n^* + D_\delta) - n^* n^*) \\
&= \frac{nn^* + n^* D_\delta}{nn^* + nD_\delta} \leq 1 \ .
\end{aligned}$$

□

11.3. *Our Allocation Function is 1-Distinguishably Monotonic*

LEMMA 11.3 *$f^{(\delta)}$ is monotonic.*

Proof. Since $f^{(\delta)}$ is symmetric, it suffices to show its monotonicity with respect to a single coordinate, and we choose the n-th one for notational convenience.

Let $z_{-n} = (z_1, \ldots, z_{n-1}) \in [0, B]^{n-1}$ and assume, with no generality loss, that $z_1 \geq z_2 \geq \cdots \geq z_{n-1}$.

We must prove that if $0 \leq z_n^\perp < z_n^\top \leq B$, then

$$f_n^{(\delta)}(z_{-n}, z_n^\perp) \leq f_n^{(\delta)}(z_{-n}, z_n^\top) \ . \tag{11.8}$$

To prove Eq. 11.8 we establish two convenient claims. Before doing so, we wish to stress that, although z_{-n} is assumed to be monotonically non increasing, when z_n is chosen arbitrarily the profile (z_{-n}, z_n) may not be monotonic.

CLAIM 1. If n' is the number of candidate winners when only the first $n-1$ players are bidding and their bid profile is z_{-n}, then for all z_n

$$z_n \leq \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \;\Rightarrow\; f_n^{(\delta)}(z_{-n}, z_n) = 0 \;\; \textit{(i.e., } n \textit{ is a loser)} \tag{11.9}$$

$$z_n > \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \;\Rightarrow\; f_n^{(\delta)}(z_{-n}, z_n) > 0 \;\; \textit{(i.e., } n \textit{ is a winner)} \tag{11.10}$$

Proof of CLAIM 1. The hypothesis of Claim 1 can be re-written as follows:

$$\forall i \in \{1, 2, \ldots, n'\},\; z_i > \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta};\; \text{and } \forall i \in \{n'+1, \ldots, n-1\},\; z_i \leq \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \;\; .$$

Let ("player n join the game bidding") $z_n \leq \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$. Then

$$\forall i \in \{1, 2, \ldots, n'\}, z_i > \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta};\; \text{and } \forall i \in \{n'+1, \ldots, n\}, z_i \leq \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \;\; .$$

That is, the bid threshold continues to be $\frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$, and the set of winners continues to be $\{1, 2, \ldots, n'\}$. Thus n is a loser and inequality 11.9 holds.

Let now (player n join the game bidding) $z_n > \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$ and assume, for the sake of contradiction, that $f_n^{(\delta)}(z_{-n}, z_n) = 0$, that is, that player n is a loser. Then, letting n^* be the new number of candidate winners, by definition:

$$\forall i \in \{1, 2, \ldots, n^*\}, z_i > \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta}; \quad \forall i \in \{n^*+1, \ldots, n\}, z_i \leq \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta} \;\; .$$

Thus, "ignoring n" we get

$$\forall i \in \{1, 2, \ldots, n^*\}, z_i > \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta}; \quad \forall i \in \{n^*+1, \ldots, n-1\}, z_i \leq \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta} \;\; .$$

That is, n^* is also the number of candidate winners under the hypothesis of Claim 1. Thus, the uniqueness of n^* implies $n^* = n'$. In turn, this implies that $z_n \leq \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$, a contradiction. This contradiction proves that n is a winner. Thus the claimed inequality (11.10) holds. *Q.E.D.*

Thanks to Claim 1, to establish the monotonicity of $f^{(\delta)}$ it suffices to prove that

$$(11.11)\quad \textbf{if } \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} < z_n^\perp < z_n^\top, \textbf{ then } f_n^{(\delta)}(z_{-n}, z_n^\perp) \le f_n^{(\delta)}(z_{-n}, z_n^\top) \ .$$

Notice that for such $z_n^\perp$ and $z_n^\top$, player n is always a candidate winner. Therefore, let $\{1, \ldots, n^\perp, n\}$ and $\{1, \ldots, n^\top, n\}$ be the winners when the bid profiles are $(z_{-n}, z_n^\perp)$ and $(z_{-n}, z_n^\top)$ respectively. We now relate $n^\perp$ and $n^\top$.

CLAIM 2. $\quad n^\perp \ge n^\top$.

Proof of CLAIM 2. Assume by way of contradiction that $n^\perp < n^\top$.
We proceed in a way similar to the proof of Lemma 11.2. Set

$$n^\Delta \stackrel{\text{def}}{=} n^\top - n^\perp, \quad S^\perp \stackrel{\text{def}}{=} \sum_{j=1}^{n^\perp} z_j, \quad \text{and} \quad S^\top \stackrel{\text{def}}{=} \sum_{j=1}^{n^\top} z_j = S^\perp + S^\Delta.$$

Since $n^\perp \le i < n^\top$ implies that player i is a loser when the bid profile is $(z_{-n}, z_n^\perp)$ and a winner when the bid profile is $(z_{-n}, z_n^\top)$, we have

$$\frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} \ge z_i > \frac{S^\top + z_n^\top}{n^\top + 1 + D_\delta} = \frac{S^\perp + S^\Delta + z_n^\top}{n^\perp + n^\Delta + 1 + D_\delta} \ .$$

Averaging over all i such that $n^\perp \le i < n^\top$ we get:

$$(11.12)\quad \frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} \ge \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + S^\Delta + z_n^\top}{n^\perp + n^\Delta + 1 + D_\delta} \ .$$

Focusing on the second inequality of (11.12), we have

$$\begin{aligned} \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + S^\Delta + z_n^\top}{n^\perp + n^\Delta + 1 + D_\delta} &\Leftrightarrow (n^\perp + n^\Delta + 1 + D_\delta)S^\Delta > n^\Delta(S^\perp + S^\Delta + z_n^\top) \\ &\Leftrightarrow (n^\perp + 1 + D_\delta)S^\Delta > n^\Delta(S^\perp + z_n^\top) \\ &\Leftrightarrow \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + z_n^\top}{n^\perp + 1 + D_\delta} \ . \end{aligned}$$

Thus, since $z_n^\perp < z_n^\top$, the second inequality (11.12) contradicts the first.

The contradiction establishes that $n^\perp \ge n^\top$ as claimed. *Q.E.D.*

We now use the fact that $n^\perp \geq n^\top$ to prove Eq. 11.11, as desired.

If $n^\perp = n^\top$, then, for both $(z_{-n}, z_n^\top)$ and $(z_{-n}, z_n^\perp)$, the set of candidate winners is $\{1, 2, \ldots, n^\perp, n\}$. Thus, letting $n^* = n^\perp + 1 = n^\top + 1$ be the number of candidate winners, we get

$$\begin{aligned} f_n^{(\delta)}(z_{-n}, z_n^\perp) &= \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} \cdot \frac{z_n^\perp(n^* + D_\delta) - \sum_{j=1}^{n^*-1} z_j - z_n^\perp}{z_n^\perp D_\delta} \\ &\leq \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} \cdot \frac{z_n^\top(n^* + D_\delta) - \sum_{j=1}^{n^*-1} z_j - z_n^\top}{z_n^\top D_\delta} = f_n^{(\delta)}(z_{-n}, z_n^\top) \ . \end{aligned}$$

If $n^\perp > n^\top$, then let $n^\perp = n^\top + n^\Delta$, $S^\top = \sum_{j=1}^{n^\top} z_j$ and $S^\perp = \sum_{j=1}^{n^\perp} z_j = S^\top + S^\Delta$ as before. Averaging over all z_i such that $n^\top < i \leq n^\perp$ we get:

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} = \frac{S^\top + S^\Delta + z_n^\perp}{n^\top + n^\Delta + 1 + D_\delta} \ . \tag{11.13}$$

But this is equivalent to

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\top + z_n^\perp}{n^\top + 1 + D_\delta} \ . \tag{11.14}$$

Letting $C_1 = \frac{n+D_\delta}{n}$ and $C_2 = C_1 \frac{1}{(n^\perp+1+D_\delta)z_n^\perp} \frac{1}{(n^\top+1+D_\delta)z_n^\top}$, we now do the final calculation:

$$\begin{aligned} & f_n^{(\delta)}(z_{-n}, z_n^\top) - f_n^{(\delta)}(z_{-n}, z_n^\perp) \\ &= C_1 \cdot \Big(\frac{z_n^\top(n^\top + 1 + D_\delta) - S^\top - z_n^\top}{(n^\top + 1 + D_\delta) z_n^\top} - \frac{z_n^\perp(n^\perp + 1 + D_\delta) - S^\perp - z_n^\perp}{(n^\perp + 1 + D_\delta) z_n^\perp} \Big) \\ &= C_1 \cdot \Big(\frac{S^\perp + z_n^\perp}{(n^\perp + 1 + D_\delta) z_n^\perp} - \frac{S^\top + z_n^\top}{(n^\top + 1 + D_\delta) z_n^\top} \Big) \\ &= C_2 \cdot \Big((S^\perp + z_n^\perp)(n^\top + 1 + D_\delta) z_n^\top - (S^\top + z_n^\top)(n^\perp + 1 + D_\delta) z_n^\perp \Big) \\ &= C_2 \cdot \Big((S^\top + S^\Delta + z_n^\perp)(n^\top + 1 + D_\delta) z_n^\top - (S^\top + z_n^\top)(n^\top + n^\Delta + 1 + D_\delta) z_n^\perp \Big) \\ &= C_2 \cdot \Big(S^\top(n^\top + 1 + D_\delta)(z_n^\top - z_n^\perp) + S^\Delta(n^\top + 1 + D_\delta) z_n^\top - n^\Delta(S^\top + z_n^\top) z_n^\perp \Big) \\ &\geq C_2 \cdot \Big(S^\top(n^\top + 1 + D_\delta)(z_n^\top - z_n^\perp) + S^\Delta(n^\top + 1 + D_\delta) z_n^\top - n^\Delta(S^\top + z_n^\perp) z_n^\top \Big) \geq 0 \end{aligned}$$

The last inequality has been derived using the fact that $z_n^\top - z_n^\perp \geq 0$ and (by Eq. 11.14) the fact that $S^\Delta(n^\top + 1 + D_\delta) - n^\Delta(S^\top + z_n^\perp) > 0$.

This finishes the proof that $f^{(\delta)}$ is monotonic. $\square$

LEMMA 11.4 *$f^{(\delta)}$ is 1-distinguishably monotonic.*

Proof. We already know from Lemma 11.3 that $f^{(\delta)}$ is monotonic.

We now need to argue the integrability of $f^{(\delta)}$, that is, that $f_i^{(\delta)}(z_i, z_{-i})$ is integrable in z_i for any choice of z_{-i}. Again, because $f^{(\delta)}$ is symmetric, it suffices to argue this for $i = n$. To do so, let us first point out that the number of winners does not increase when z_n increases, for any z_{-n}. Indeed, letting n' be the number of candidate winners when only the first $n-1$ players are bidding and their bid profile is z_{-n}, CLAIM 1 in the proof Lemma 11.3 tells us that when $z_n \leq \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$ player n is always a loser and thus the number of winners remains constant; moreover, CLAIM 2 in the proof of Lemma 11.3 tells us that when z_n increases past $\frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$ the number of winners does not decrease. Thus, $f^{(\delta)}$ is piecewise continuous (as we have established that the number of winners does not increase when z_n increases, the finite number of continuous pieces is at most n), and therefore $f^{(\delta)}$ is integrable.

We are therefore left to prove the "distinguishability condition".

Fix a player $i \in [n]$ and two distinct valuations $v_i, v_i' \in \{0, 1, \ldots, B\}$, and assume that $v_i < v_i'$. If we choose $v_{-i} = (v_i, v_i, \ldots, v_i)$, then:

- $f_i^{(\delta)}(v_i, v_{-i}) = \frac{1}{n}$ since there are n winners, all bidding the same valuation, and
- $f_i^{(\delta)}(z, v_{-i}) = \frac{1}{nD_\delta}(D_\delta + n - 1 - \frac{v_i}{z}(n-1)) > \frac{1}{n}$, when $v_i < z \leq (1+D_\delta)v_i$.

The upper bound $(1 + D_\delta)v_i$ is to make sure that the number of winners is still n on input (z, v_{-i}). Thus, by definition, $f_i^{(\delta)}(z, v_{-i})$ is a function that is *strictly* increasing when z increases in $(v_i, (1 + D_\delta)v_i]$, and therefore

$$\int_{v_i}^{v_i'} \left(f_i^{(\delta)}(z, v_{-i}) - f_i^{(\delta)}(v_i, v_{-i})\right) dz \geq \int_{v_i}^{\min\{v_i', (1+D_\delta)v_i\}} \left(f_i^{(\delta)}(z, v_{-i}) - f_i^{(\delta)}(v_i, v_{-i})\right) dz > 0 \ ,$$

as desired. □

11.4. *Our Allocation Function is δ-Good*

Our last lemma tells us that $M_{\mathsf{opt}}^{(\delta)} \stackrel{\text{def}}{=} M_{f^{(\delta)}}$ is a 1-distinguishably monotone mechanism. This property (via our Distinguishable Monotonicity Lemma) simplifies the analysis of the undominated strategies for $M_{\mathsf{opt}}^{(\delta)}$, but otherwise has no bearing on proving the social-welfare performance claimed for $M_{\mathsf{opt}}^{(\delta)}$ in Theorem 4. (As already remarked, the probabilistic second-price mechanism is 1-DM, but only guarantees a fraction $\left(\frac{1-\delta}{1+\delta}\right)^2$ of the maximum social welfare.) Accordingly, our allocation function $f^{(\delta)}$ must and indeed does satisfy an additional property, *δ-goodness*. We state this property below and prove that $f^{(\delta)}$ indeed satisfies it. Only in the next section we shall prove the relevance of δ-goodness for Theorem 4.

Recall that $D_\delta \stackrel{\text{def}}{=} \left(\frac{1+\delta}{1-\delta}\right)^2 - 1$.

DEFINITION 11.5 *We say that a 1-DM allocation function f is* **δ-good** *if*

$$\forall\, i \in \{1,\ldots,n\},\ \forall\, v \in \{0,1,\ldots,B\}^n, \quad \sum_{j=1}^{n} f_j(v)v_j + D_\delta \cdot f_i(v)v_i \geq \frac{1}{n}\cdot v_i(n+D_\delta)\ .$$

LEMMA 11.6 *$f^{(\delta)}$ is δ-good.*

PROOF: As we already know that $f^{(\delta)}$ is 1-DM, we establish the lemma by proving that the above inequality holds not only for the discrete cube $\{0,1,\ldots,B\}^n$, but actaually for the continuous cube $[0,B]^n$.

Without loss of generality, assume $z_1 \geq z_2 \geq \cdots \geq z_n$. Observe that:

$$\begin{aligned}\sum_{i=1}^{n} f_i^{(\delta)}(z)z_i = \sum_{i=1}^{n^*} f_i^{(\delta)}(z)z_i &= \frac{1}{n}\cdot\frac{n+D_\delta}{n^*+D_\delta}\cdot\sum_{i=1}^{n^*}\frac{z_i(n^*+D_\delta)-\sum_{j=1}^{n^*} z_j}{D_\delta} \\ &= \frac{1}{n}\cdot\frac{n+D_\delta}{n^*+D_\delta}\cdot\left(\sum_{i=1}^{n^*} z_i\right)\ .\end{aligned}$$

Then, for each looser k (i.e., each player $k > n^*$, we have

$$\sum_{j=1}^{n} f_j^{(\delta)}(z)z_j + D_\delta\cdot f_k^{(\delta)}(z)z_k = \sum_{i=1}^{n} f_i^{(\delta)}(z)z_i = \frac{n+D_\delta}{n}\cdot\frac{\sum_{i=1}^{n^*} z_i}{n^*+D_\delta} \geq \frac{n+D_\delta}{n}\cdot z_k$$

where the last inequality is due to Eq. 11.2.

At the same time, for each winner i (i.e., each player $i \le n^*$), we have

$$\sum_{j=1}^{n} f_j^{(\delta)}(z) z_j + D_\delta \cdot f_i^{(\delta)}(z) z_i = \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} \cdot \left(\sum_{i=1}^{n^*} z_i \right) + D_\delta \cdot f_i^{(\delta)}(z) z_i$$
$$= \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} z_i (n^* + D_\delta) = \frac{1}{n} \cdot z_i (n + D_\delta)$$

again satisfying Eq. 11.5. *Q.E.D.*

11.5. *Proof of Theorem 4*

Since $M_{\mathsf{opt}}^{(\delta)} = M_{f^{(\delta)}}$, Theorem 4 immediately follows from the fact that our allocation function $f^{(\delta)}$ is δ-good and the following

LEMMA 11.7 *If f is δ-good, then for the mechanism $M_f = (S, F)$ we have: for all contexts $(n, B, \delta, K, \theta)$ and all strategy profile $v \in \mathsf{UDed}(K)$,*

$$\mathbb{E}\Big[\mathrm{SW}\big(\theta, F(v)\big)\Big] \ge \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} \right) \mathrm{MSW}(\theta) \ .$$

PROOF: Arbitrarily fix a context $(n, B, \delta, K, \theta)$ and $v \in \mathsf{UDed}(K)$. Then, because in any allocation the social welfare coincides with the true valuation of some player, to prove our lemma it suffices to prove that

$$(11.15) \quad \forall i \in [n], \quad \sum_{j=1}^{n} \theta_j f_j(v) \ge \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} \right) \theta_i \ .$$

For every $i \in [n]$, let $x_i \in \mathbb{R}$ be such that $K_i \subseteq \delta[x_i]$, and let $\delta[x] = \delta[x_1] \times \cdots \times \delta[x_n]$. Then, the Distinguishable Monotonicity Lemma respectively implies $(1-\delta)x_i \le \min K_i \le v_i \le \max K_i \le (1+\delta)x_i$; equivalently,

$$\frac{v_i}{1+\delta} \le x_i \le \frac{v_i}{1-\delta} \ .$$

Also, $\theta \in K$ implies $(1-\delta)x_i \le \theta_i \le (1+\delta)x_i$. Combining the last two chains of inequalities yields

$$(11.16) \quad \frac{1-\delta}{1+\delta} v_i \le \theta_i \le \frac{1+\delta}{1-\delta} v_i \ .$$

Let us now argue that Eq. 11.15 holds by arbitrarily fixing v and i and showing that it is impossible to construct a "bad" θ so as to violate Eq. 11.15.

In trying to construct a "bad" θ, it suffices to choose θ_j (for $j \neq i$) to be as small as possible, since θ_j only appears on the left-hand side with a positive coefficient. For θ_i, however, we may want to choose it as large as possible if $f_i(v) \geq \left(\frac{(1-\delta)^2+\frac{4\delta}{n}}{(1+\delta)^2}\right)$, or as small as possible otherwise. So there are two extreme θ's.

Considering these extreme choices, we conclude that no θ contradicts Eq. 11.15 if:

$$\sum_{j=1}^{n} \left(\frac{1-\delta}{1+\delta}\right) v_j f_j(v) \geq \left(\frac{(1-\delta)^2+\frac{4\delta}{n}}{(1+\delta)^2}\right) \left(\frac{1-\delta}{1+\delta}\right) v_i \ , \text{ and}$$

$$\sum_{j=1}^{n} \left(\frac{1-\delta}{1+\delta}\right) v_j f_j(v) + \left(\frac{1+\delta}{1-\delta} - \frac{1-\delta}{1+\delta}\right) v_i f_i(v) \geq \left(\frac{(1-\delta)^2+\frac{4\delta}{n}}{(1+\delta)^2}\right) \left(\frac{1+\delta}{1-\delta}\right) v_i \ .$$

Simplifying the above equations, Eq. 11.15 holds if both the following inequalities hold:

$$(11.17) \quad \sum_{j=1}^{n} v_j f_j(v) \geq \frac{n+D_\delta}{n} \cdot \frac{1}{D_\delta+1} \cdot v_i \ ,$$

$$(11.18) \quad \sum_{j=1}^{n} v_j f_j(v) + D_\delta \cdot v_i f_i(v) \geq \frac{n+D_\delta}{n} v_i \ .$$

Note that Eq. 11.18 holds because it is implied by the hypothesis that f is δ-good; note also that Eq. 11.17 holds because it is implied by Eq. 11.18. Indeed, since $\frac{1}{D_\delta+1} = \left(\frac{1-\delta}{1+\delta}\right)^2 < 1$ for all $\delta \in (0,1)$,

$$\sum_{j=1}^{n} v_j f_j(v) \geq \frac{1}{D_\delta+1} \left(\sum_{j=1}^{n} v_j f_j(v) + D_\delta v_i f_i(v) \right) \geq \frac{1}{D_\delta+1} \frac{n+D_\delta}{n} v_i \ .$$

Thus both Eq. 11.15 and both our lemma and Theorem 4 hold. *Q.E.D.*

11.6. *The Computational Efficiency of* $M_{\mathsf{opt}}^{(\delta)}$

Finally, we wish to clarify that, although $M_{\mathsf{opt}}^{(\delta)} = M_{f^{(\delta)}} = (S, F)$ is not as simple as the second-price mechanism, it can still be efficiently implemented. That is, both the allocation function $F^A = f^{(\delta)}$ and the expected price function F^P are efficiently computable over $\{0, 1, \ldots, B\}^n$.

The computational efficiency of F^A is apparent once one realizes that the number of candidate winners, n^*, can be determined in linear time.

The computational efficiency of F^P requires a bit of an argument. Without loss of generality, let us show how to compute the expected price for player n. Recall that, for a bid profile v,

$$F_n^P(v_{-n}, v_n) = f_n^{(\delta)}(v_{-n}, v_n) \cdot v_n - \int_0^{v_n} f_n^{(\delta)}(v_{-n}, z)\, dz \ .$$

When v_{-n} is fixed, $f_n^{(\delta)}$ is a function piece-wisely defined according to v_n, since different values of v_n may result in different numbers of winners n^*. Assume without loss of generality that $v_1 \geq v_2 \geq \cdots \geq v_{n-1}$, and let n' be the number of winners when player n is absent.

When $v_n \leq \frac{\sum_{j=1}^{n'} v_j}{n'+D_\delta}$, the proof of the monotonicity of $f^{(\delta)}$ implies that $f_n^{(\delta)} = 0$, so that integral below this point is zero.

When $v_n > \frac{\sum_{j=1}^{n'} v_j}{n'+D_\delta}$, one can again see from the proof of the monotonicity of $f^{(\delta)}$ that n^* is non-increasing as a function of v_n. Therefore, $f_n^{(\delta)}$ contains at most n different pieces and, for each piece with n^* fixed, $f_n^{(\delta)}(v_{-n}, v_n) = a + b/v_n$ is a function that is symbolically integrable. Therefore, the only question is how to calculate the pieces for $f_n^{(\delta)}$.

Conceptually, one starts from $v_n = \frac{\sum_{j=1}^{n'} v_j}{n'+D_\delta}$ and "moves v_n upwards", recording the points at which Eq. 11.2 is violated, because these are the "borders of the continuous pieces" of $f_n^{(\delta)}$. Practically, this seemingly infinite procedure may be efficiently carried out by a *line sweep method*.

12. CONCLUSIONS

Mechanism design is undoubtedly a fascinating field. One can only marvel at the possibility that an ignorant social planner can leverage the knowledge and the rationality of the players in order to obtain the outcomes he desires. But if we want to transform this beautiful theory into strong guarantees in the real world it is important to minimize its underlying assumptions.

To us, the assumption that each player knows "on the nose" his own valuation appears to be too idealized in many an environment. Even the assumption that each player knows a probability distribution from which his own true valuation has been drawn is very strong. To be safe, we should budget for the possibility of Knightian players.

In any field, as we progress from idealized to more and more realistic models, we should expect to face additional complexities. Knightian mechanism design will be no exception. Nonetheless, we should remain optimistic. At least for single-good auctions, Knightian mechanism design is workable.

ACKNOWLEDGEMENTS

We are grateful to Alan Deckelbaum for pointing out simplifications to our proofs.

APPENDIX A: PERFORMANCE DIAGRAMS

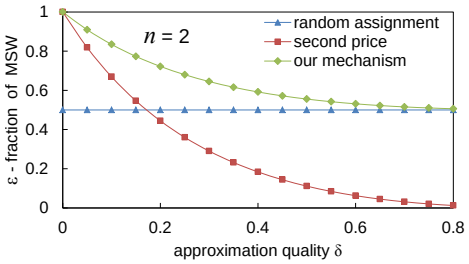

(a) With $n = 2$ players, the second-price mechanism performs *worse* than randomly assigning the good for $\delta > 0.18$.

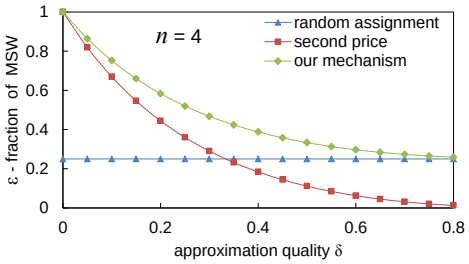

(b) With $n = 4$ players, the second-price mechanism performs *worse* than randomly assigning the good for $\delta > 0.34$.

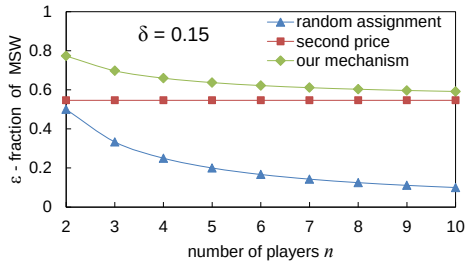

(c) With $\delta = 0.15$, the second-price mechanism always performs better than randomly assigning the good.

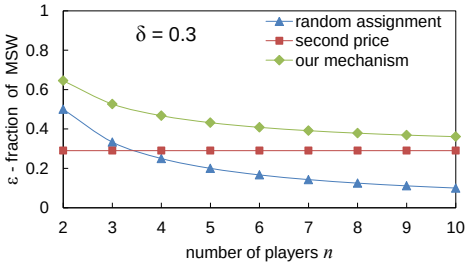

(d) With $\delta = 0.3$, the second-price mechanism performs *worse* than randomly assigning the good for $n = 2, 3$.

Figure 1: The social welfare guarantees of randomly assigning the good ($\varepsilon = \frac{1}{n}$), the second-price mechanism ($\varepsilon = \frac{(1-\delta)^2}{(1+\delta)^2}$), and our optimal mechanism ($\varepsilon = \frac{(1-\delta)^2+\frac{4\delta}{n}}{(1+\delta)^2}$). In (1a) and (1b) we compare ε versus δ, and in (1c) and (1d) we compare ε versus n. The green data, our mechanism, is always better (at times significantly) than the other two mechanisms.

REFERENCES

Robert J. Aumann. Utility theory without the completeness axiom. *Econometrica*, 30 (3):445–462, July 1962.

Truman F. Bewley. Knightian decision theory. Part I. *Decisions in Economics and Finance*, 25(2):79–110, 2002. Earlier version appeared as a discussion paper no. 807 of the Cowles Foundation at Yale University, November 1986.

Eric Danan. Randomization vs. selection: How to choose in the absence of preference? *Management Science*, 56:503–518, March 2010.

Juan Dubra, Fabio Maccheroni, and Efe A. Ok. Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115(1):118–133, March 2004.

Uriel Feige and Moshe Tennenholtz. Mechanism design with uncertain inputs: (to err is human, to forgive divine). In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, STOC '11, pages 549–558, New York, NY, USA, 2011. ACM.

Vincy Fon and Yoshihiko Otani. Classical welfare theorems with non-transitive and non-complete preferences. *Journal of Economic Theory*, 20(3):409–418, June 1979.

D. Gale and A. Mas-Colell. An equilibrium existence theorem for a general model without ordered preferences. *Journal of Mathematical Economics*, 2(1):9–15, March 1975.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, April 1989.

Matthew O Jackson. Implementation in undominated.strategies: A look at bounded mechanisms. *Review of Economic Studies*, 59(4):757–75, October 1992.

Frank H. Knight. *Risk, Uncertainty and Profit*. Houghton Mifflin, 1921.

Giuseppe Lopomo, Luca Rigatti, and Chris Shannon. Uncertainty in mechanism design, 2009. `http://www.pitt.edu/~luca/Papers/mechanismdesign.pdf`.

Andrew Mas-Colell. An equilibrium existence theorem without complete or transitive preferences. *Journal of Mathematical Economics*, 1(3):237–246, December 1974.

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

Leandro Nascimento. Remarks on the consumer problem under incomplete preferences. *Theory and Decision*, 70(1):95–110, January 2011.

Efe A. Ok. Utility representation of an incomplete preference relation. *Journal of Economic Theory*, 104:429–449, 2002.

Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172:1783–1799, October 2008.

Luca Rigotti and Chris Shannon. Uncertainty and risk in financial markets. *Econometrica*, 73(1):203–243, 01 2005.

Tuomas Sandholm. Issues in computational vickrey auctions. *International Journal of Electronic Commerce*, 4:107–129, March 2000.

David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–87, May 1989.

Wayne Shafer and Hugo Sonnenschein. Equilibrium in abstract economies without ordered preferences. *Journal of Mathematical Economics*, 2(3):345–348, December 1975.