Mechanism Design with Approximate Valuations

Alessandro Chiesa* MIT CSAIL

Silvio Micali† MIT CSAIL

Zeyuan Allen Zhu‡ MIT CSAIL

December 7, 2011

Abstract

We study single-good, incomplete-information auctions in a Knightian setting that is *approximate*: that is, when each player knows his true payoff type only within a constant factor $\delta \in (0, 1)$. On the negative side, we prove that in our setting no dominant-strategy mechanism can significantly guarantee better social welfare 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.

Pay attention that there are some different versions of this paper:

- This version is maintained by the third author, and consistent with his Master Thesis. It replaces two earlier versions with minor mistakes appeared as technical reports:
  - MIT-CSAIL-TR-2011-009, MIT, February 2011, and
  - MIT-CSAIL-TR-2011-024, MIT, April 2011.
- A parallel version maintained by the first two authors with slightly different notions and proofs appeared as technical report:
  - `http://arxiv.org/abs/1112.1147` on arxiv, entitled "Knightian Auctions".

*Email: alexch@csail.mit.edu.

†Email: silvio@csail.mit.edu.

‡Email: zeyuan@csail.mit.edu.

Contents

1 Brief Introduction

A general direction and a specific focus. Mechanism design aims at producing a desired outcome by leveraging each player's rationality and his knowledge of his own (pay-off) type. But:

What happens when each player knows his own type only **approximately***?*

We explore this general direction focusing on a specific goal (trivially achieved when the players know their own types exactly):

Guaranteeing *high social welfare in auctions of a single good.*

We do so from a *finite* perspective: namely, we consider players with only finitely many possible types and mechanisms specifying for each player only finitely many strategies.

Our model of self uncertainty. In a single-good auction, a player type is a natural number, referred to as a *valuation.* The possibility that a player in such an auction may not precisely know his own valuation strikes us to be quite realistic. No one would be too surprised if, tasked to figure out a firm's *true* valuation for the good, different employees reported different values, or some of them reported ranges of values rather than single values. This said, for decision theory —let alone mechanism design!— to be meaningful at all, the players must know *something* about their own types. The classical work of [Kni21] (see also [Bew02]) envisages that each player knows that his own type is distributed according to one of several possible distributions. Other works —e.g., [Mye81, San00, FT11]— envisage a more structured kind of "self knowledge": namely, each player knows *the* distribution from which his own type has been drawn.

Our model of the players' "self uncertainty" is purely *set-theoretic.* In essence, each player only knows that his own valuation is one of several candidates. In this model we investigate how good mechanisms one can design when it is the case that the candidates of each player are *relatively clustered:* that is, when "each player knows his own valuation within the same fixed percentage".

Roadmap. In the main body of this extended abstract we motivate our model and its associated solution concepts, state our three main results, highlight two techniques that we believe to be of independent interest, and finally compare our model and work with prior ones. All formalizations and proofs can be found in the appendix.

2 The Approximate-Valuation Model

A δ*-approximate context*, where δ is a constant in $[0, 1)$, is formally defined in Appendix A. Informally, in such a context all possible valuations are integers between 0 and a *valuation bound* B. Each player i does not know his own *true valuation*, θ_i, but a set $K_i \subseteq \{0, \ldots, B\}$ such that

$$\text{(i) } \theta_i \in K_i \qquad \text{and} \qquad \text{(ii) } K_i \subseteq \delta[c_i],$$

where c_i is the "center" of K_i and, for all $x \in \mathbb{R}$, $\delta[x]$ consists of all possible valuations within $x \pm \delta x$, that is, $\delta[x] \stackrel{\text{def}}{=} \big[(1-\delta)x, (1+\delta)x\big] \cap \{0, 1, \ldots, B\}$.

We refer to δ as an *approximation accuracy*, and to each K_i as the *approximate-valuation set* of player i (or δ*-approximate-valuation set* if we wish to be more precise). We denote by $\mathscr{C}^{\delta}_{n,B}$ the class of *all* δ-approximate contexts with n players and valuation bound B.

Each approximate-valuation set K_i should be interpreted as the set of *all and only* candidates for θ_i in i's mind. For instance, when $\delta = 0$ he knows his own valuation *exactly*, and when $\delta = 0.1$ "within a 10% accuracy". No matter how accurately each player knows his own true valuation, every context is δ-approximate for a suitably large δ: after all, all contexts are 1-approximate!

Q&A.

- *Why δ?* In our purely set-theoretic framework, the approximation accuracy δ can be interpreted as quantifying the "quality of the players' knowledge about themselves". We thus find it natural to measure the performance of a mechanism as a function of δ. Without identifying any structure in the players' possible multiple valuations, one may at most design elementary mechanisms, rather than "good" single-good auctions. In sum, the accuracy parameter δ is our Trojan horse for bringing meaningful mechanism design into our model.
- *Why $\theta_i \in K_i$?* A player's approximate-valuation set contains the player's true valuation because we consider each player to be "the ultimate authority about himself".
- *Why is K_i not just an interval of integers?* An approximate-valuation set K_i may indeed be just that. But defining it to be a more general subset of $\{0, \ldots, B\}$ may be necessary in some contexts. For instance, consider a player i who is about to participate to a yard-sale auction of a large sofa. He may know precisely the amount v he would pay for the sofa, but also that, if he wins it, he would have to carry it on top of his car, which is illegal and punishable with a fine f. In a Bayesian setting, he should compute his expected value of the sofa from the probability of being caught by the police on his way home (presumably based on the specific time of day, the immediate weather forecast, the immediate traffic forecast, the likelihood that other crimes might compete for the police's attention, and so on). But in our set-theoretic model, his approximate-valuation set K_i blissfully consists of just two separate values: namely, $K_i = \{v, v - f\}$.[1]
- *Multiplicative or additive accuracy?* A greater level of generality is achieved by considering two distinct accuracy parameters: a multiplicative one, δ^*, and an additive one, δ^+, leading to the following modified constraint:

  $$K_i \subseteq \left[(1-\delta^*)x_i - \delta^+, (1+\delta^*)x_i + \delta^+\right] \cap \{0, 1, \ldots, B\} \text{ for some value } x_i \in \mathbb{R}.$$

  All of our theorems hold for such more general approximate valuations. For simplicity, however, we consider only one kind of accuracy parameters, and we find the multiplicative one more meaningful.
- *Multiple possible δ's?* Yes: indeed, $\delta > \delta'$ implies that every δ'-approximate context also is a δ-approximate one, and that $\mathscr{C}^{\delta}_{n,B} \supseteq \mathscr{C}^{\delta'}_{n,B}$.
- *Do the players know δ?* A player i in a δ-approximate context may know nothing about a "global δ". Of course, knowing K_i, he can certainly compute his smallest "local" δ_i: namely, $\frac{\max K_i - \min K_i}{\max K_i + \min K_i}$. But he may not have enough information about his opponents to realize that he is in a δ-approximate context for $\delta < 1$.
- *Does the designer know δ?* When disproving the existence of mechanisms with a given efficiency guarantee for $\mathscr{C}^{\delta}_{n,B}$, we gladly assume that the mechanism designer does know δ precisely, since this makes our impossibility results stronger. When proving the existence of such mechanisms, we shall specify whether or not the designer knows a "sufficient" δ.
- *Can real δ's be really large?* Absolutely. Valuations may indeed be "very approximate". Consider a firm participating to an auction for the exclusive rights 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 were to come from the firm's ability to lower its costs of production via some breakthrough research, an approximation accuracy of the firm's own valuation for the license could easily exceed 0.5.

[1] Note that such fine it is not paid to the seller, and cannot be modeled "within the game". It is an element extraneous to the auction, but clearly affecting the valuation of our particular player.

3 Appropriate Solutions Concepts and Performance Measures

Solution concepts. In a non-Bayesian setting of incomplete information, two notions of implementations are natural to explore: implementation in dominant strategies and in undominated strategies [Jac92]. However, the classical definitions of *dominance* need to be extended in order to apply to our approximate-valuation model. (Essentially, we adapt the generalized notions of dominance for Knightian uncertainty [Kni21, Bew02] to our purely set-theoretic setting.)

Informally, a strategy s_i of a player i *very-weakly dominates* another strategy t_i of i, relative to a mechanism M and an approximate-valuation set K_i of i, if, for all candidate valuations in K_i and all possible strategy subprofiles of his opponents, s_i gives i a utility greater than or equal to that given him by t_i. If it is further the case that s_i gives i strictly greater utility than t_i for at least some valuations in K_i and strategy subprofiles of i's opponents, then s_i *weakly dominates* t_i. A strategy of i is *undominated* if it is not weakly dominated. (The set of such undominated strategies is denoted by $\mathsf{UDed}_i^M(K_i)$, or simply by $\mathsf{UDed}_i(K_i)$ when M is clear from context.)

Performance measures. As mentioned at the beginning, our plan is to provide "worst-case guarantees" about social welfare for single-good auctions in the approximate-valuation setting. This plan requires some explaining: indeed, when all knowledge resides with the players and they are uncertain about their own valuations, what should "maximum social welfare" and "actual social welfare" mean? Conceptually, we envisage the following process:

1. A context in $\mathscr{C}_{n,B}^{\delta}$ materializes: that is, there is one good for sale and n players show up, each player i with a δ-approximate-valuation set set K_i.
2. A mechanism designer, knowing only n and B (and in some applications also a valid accuracy parameter δ), chooses a solution concept and constructs a (possibly probabilistic) mechanism M for auctioning the good.
3. The "devil", knowing everything specified so far, secretly selects a true valuation profile θ such that $\theta_i \in K_i$ for every player i.
4. Each player i, based on his approximate-valuation set K_i, selects (possibly probabilistically) and reports a strategy s_i in the set of strategies S_i provided to him by M. (Perhaps, player i may learn θ_i after the auction is over. Perhaps, he may never learn it.)
5. The mechanism then evaluates its (possibly probabilistic) outcome function F on the reported strategy profile s so as to produce an outcome $\omega = (j, P)$: that is, the outcome $F(s)$ specifies the player j winning the good, and the profile of prices $P = (P_1, \ldots, P_n)$ that the players pay.

Given this process, whether or not the players eventually become aware of their own true valuations, the maximum social welfare relative to the secret devil-chosen θ, $\mathrm{MSW}(\theta)$, is taken to be $\max_i \theta_i$, and the actual social welfare relative to θ for outcome $w = (j, P)$, $\mathrm{SW}(\theta, \omega)$ is taken to be θ_j. We are interested in studying, as a function of δ and the chosen solution concept, the expected value (over all possible sources of randomness) of the ratio

$$\frac{\mathrm{SW}(\theta, \omega)}{\mathrm{MSW}(\theta)} \ .$$

4 Results

How much social welfare can we guarantee in approximate-valuation auctions?

In a classical setting the answer is trivial: 100% in (very-weakly-)dominant strategies, via the second-price mechanism. The situation is quite different with approximate valuations.

4.1 The Inadequacy of Dominant-Strategy Mechanisms

A bit superficially, one might argue that very-weakly-dominant-strategy mechanisms cannot be meaningful in the approximate-valuation world as follows: *If a player has multiple possible candidates for his true valuation, how can he know which one is "the best" for him to bid no matter what his opponents do?*

This "reasoning" presupposes that, as in the exact-valuation world, an auction mechanism can safely restrict a player's strategies to (reporting) single valuations. In our setting, however, it is not only reasonable but even natural for a mechanism to allow a player to report a *set of valuations* (e.g., his own K_i). Indeed, it is easy to realize that the revelation principle [Mye81] continues to guarantee that every very-weakly-dominant-strategy mechanism for $\mathscr{C}^{\delta}_{n,B}$ has an equivalent *very-weakly-dominant-strategy-truthful* mechanism. In our approximate world, a mechanism is of the latter kind if, for each player i: (1) S_i, the strategies of i, consists of reporting an arbitrary δ-approximate set V_i of valuations, and (2) reporting his true approximate-valuation set K_i is a very-weakly-dominant strategy.

With such richer strategy sets, in principle there might be a dominant-strategy mechanism guaranteeing maximum social welfare. More realistically, in light of the approximate accuracy of the players' self knowledge, one should expect some degradation of performance to be unavoidable. For instance, one might conjecture that, in a δ-approximate context, 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 also this is too optimistic.

Theorem 1. *For all n, $\delta \in (0,1)$, $B > \frac{3-\delta}{2\delta}$, and (possibly probabilistic) very-weakly-dominant-strategy-truthful mechanism $M = (S, F)$, there exists a δ-approximate-valuation profile K and a true-valuation profile $\theta \in K$ such that*

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

(The proof of Theorem 1 can be found in Appendix C.)

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 very 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 "stupid" 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 one". In other words,

unless we are in some form of Bayesian model,
dominant strategies are **intrinsically** *linked to the* **exact** *knowledge of our own valuations.*

A conceptual contribution. The superficial reasoning at the beginning of this section may be wrong, but not the gut feeling that dominant strategies must be "wrong" for the approximate-valuation setting! Intuition, however, must be formalized. This is what Theorem 1 does. Although not very hard to prove, this theorem is conceptually important. By formally ruling out dominant-strategy mechanisms from meaningful consideration, it opens the door to alternative solution concepts: in particular, to implementation in undominated strategies. We actually believe that our approximate-valuation setting will provide a new and vital role for this classical, robust, and non-Bayesian solution concept.

4.2 The Power of Deterministic Undominated-Strategy Mechanisms

Our next theorem states that the deterministic second-price mechanism: (1) essentially guarantees a fraction $\left(\frac{1-\delta}{1+\delta}\right)^2$ of the maximum social welfare in undominated strategies, and (2) is essentially optimal among all deterministic undominated-strategy mechanisms. More formally, denoting by $\mathsf{UDed}(K)$ the Cartesian product $\mathsf{UDed}_1(K_1) \times \cdots \times \mathsf{UDed}_n(K_n)$,

Theorem 2. *The following two statements hold:*

a. Let $M_{\mathsf{2P}} = (S_{\mathsf{2P}}, F_{\mathsf{2P}})$ be the second-price mechanism with a deterministic tie-breaking rule. Then, for all n, $\delta \in [0,1)$, B, δ-approximate-valuation profiles K, true-valuation profiles $\theta \in K$, and strategy profiles $v \in \mathsf{UDed}(K)$:

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

b. Let $M = (S, F)$ be a deterministic *mechanism. Then, for all $n, \delta \in (0,1)$ and $B \geq \frac{1+\delta}{\delta}$, there exist a δ-approximate-valuation profile K, a true-valuation profile $\theta \in K$, and a strategy profile $s \in \mathsf{UDed}(K)$ such that:*

$$\mathrm{SW}\big(\theta, F(s)\big) \leq \left(\left(\frac{1-\delta}{1+\delta}\right)^2 + \frac{3}{B}\right)\mathrm{MSW}(\theta)\ .$$

(Theorem 2.a and Theorem 2.b are respectively proven in Appendix H and Appendix E.[3])

Avoiding confusion. In the exact-valuation world, the second-price mechanism achieves perfect efficiency both in dominant strategies and in undominated strategies. But in the approximate-valuation world, it is no longer a dominant-strategy one.

Easier and harder. Theorem 2.a is not hard to prove. At a very high level,

> *"It is obvious that each player i should only consider bidding a value v_i inside his own approximate-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 $\approx \left(\frac{1-\delta}{1+\delta}\right)^2$ of the maximum social welfare.* $\square$*"*

Of course, things are a bit more complex. In particular, the fact that a player i should only consider bids in K_i requires a proof (and is actually "technically" wrong as stated).

Theorem 2.b is harder to prove, as we should expect for all impossibility results. Working in undominated strategies, the revelation principle no longer applies. Thus, rather than analyzing a single mechanism (the "direct truthful" one), in principle we should consider *all* possible mechanisms, and establish that each one of them does no better than the second-price one. In particular, while the second-price mechanism has a clear and simple strategy space (namely, an integer in $\{0, 1, \ldots, B\}$), we should consider mechanisms giving the players *absolutely arbitrary* strategies: even reporting arbitrary subsets of $\{0, 1, \ldots, B\}$ would be a strong restriction! Establishing Theorem 2.b thus requires new techniques, informally discussed in Section 5, and formally provided in Appendix D.

---

[2]Breaking ties at random, the performance guarantee is only marginally better: namely, *exactly* $(\frac{1-\delta}{1+\delta})^2\mathrm{MSW}(\theta)$.

[3]To denote a strategy profile, we use "v" in the statement of Theorem 2.a to emphasize that each player in the M_{2P} indeed bids a valuation; and "s" in Theorem 2.b to emphasize that a player's strategies can be totally arbitrary.

4.3 The Greater Power of Probabilistic Undominated-Strategy Mechanisms

Our final Theorem shows that, in undominated strategies, there exists an essentially optimal probabilistic mechanism.

Theorem 3. *The following two statements hold:*

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

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

b. Let $M = (S, F)$ *be a (deterministic or probabilistic) mechanism. Then for all* $n, \delta \in (0,1)$*, and* $B \geq \frac{1+\delta}{\delta}$*, there exist a* δ*-approximate-valuation profile* K*, a strategy profile* $s \in \mathsf{UDed}(K)$*, and a true-valuation profile* $\theta \in K$ *such that*

$$\mathbb{E}\Big[\mathrm{SW}\big(\theta, F(s)\big)\Big] \leq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} + \frac{3}{B}\right) \mathrm{MSW}(\theta) \ .$$

(The proofs of Theorem 3.a and Theorem 3.b can be found in Appendix I and Appendix F respectively. Note that the mechanism $M^{(\delta)}_{\mathsf{opt}}$ of Theorem 3.a is constructed given knowledge of δ.[4])

Theoretical significance. Theorem 3 highlights a novelty of the approximate-valuation world: 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 the approximate-valuation world, probabilistic mechanisms will enjoy a provably better performance in other applications as well.

Technical difficulty. The impossibility result in Theorem 3.b is again non-trivial, but Theorem 3.a is even much harder to prove. Indeed, it is the technically hardest one in this paper.

Practicality. Despite the technical difficulty of its proof, we would like to emphasize that mechanism $M^{(\delta)}_{\mathsf{opt}}$ actually requires almost no computation from the players, and a very small amount of computation from the mechanism. In essence, it is very *practically* played.

In addition, its performance is *practically* preferable to that of the second-price mechanism. For instance, when $\delta = 0.5$, $M^{(\delta)}_{\mathsf{opt}}$ guarantees a social welfare that is at least **five** times higher that of the second-price mechanism when there are 2 players, and at least **three** times higher when there are 4 players. Even when $\delta = 0.25$, the guaranteed performance of $M^{(\delta)}_{\mathsf{opt}}$ is almost **two** times higher than that of the second-price when there are 2 players. (For a full comparison chart, see Appendix J.)

5 Two Techniques of Independent Interest

New ventures require new tools. We thus wish to highlight two techniques that we believe will prove useful to the design and analysis of mechanisms in the approximate-valuation setting.

[4] With extra pains, however, one can actually get a reasonable performance if the mechanism only knows a good upper bound on δ.

The Undominated Intersection Lemma. To prove that a given social choice function cannot be implemented in undominated strategies in the approximate-valuation model, as it is needed for Theorem 2.b and Theorem 3.b, we wish to establish some basic structural properties about undominated strategies.

For example, as an intuitive warm-up, if the strategies S_i available to each player i simply consisted of reporting single valuations, that is, if $S_i = \{0, 1, \ldots, B\}$, would it be the case that

$$\mathsf{UDed}_i(K_i) = K_i \ ? \tag{5.1}$$

If so, this would imply the following:

$$K_i \cap \widetilde{K}_i \neq \emptyset \Rightarrow \mathsf{UDed}_i(K_i) \cap \mathsf{UDed}_i(\widetilde{K}_i) \neq \emptyset \ . \tag{5.2}$$

However, relation (5.1) is in general false. Even if the strategies given to each player are valuations between 0 and B, 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.

Relation (5.2) might hold even if relation (5.1) does not. But it is unclear that it does: the set of strategies S_i's provided by a mechanism can be absolutely arbitrary, rather than $\{0, 1, \ldots, B\}$. 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 *a sufficiently simple variant of relation* (5.2) *holds for all mechanisms*, not just the ones with $S_i = \{0, 1, \ldots, B\}$. Informally,

> *For any mechanism, no matter whether it is 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.*[5]

This simple property will be powerful enough to derive all of our impossibility results for implementation in undominated strategies.

The Distinguishable Monotonicity Lemma. To prove that a given social choice function can be implemented in undominated strategies, as it is needed for Theorem 2.a and Theorem 3.a, we are happy to work with a suitable class of restricted mechanisms, using only very special strategies and allocation functions. But what should "suitable" mean?

On one hand, these restrictions should suffice for proving Theorem 2.a and Theorem 3.a. On the other hand, they should ensure that the undominated strategies corresponding to a given approximate-valuation set can be characterized in a way that is both *conceptually simple* and *easy to work with.*

Specifically, we consider mechanisms whose strategies consist of possible valuations, namely the set $\{0, \ldots, B\}$, and whose allocation functions are restrictions (to $\{0, 1, \ldots, B\}^N$) of integrable functions (over $[0, B]^N$) satisfying a suitable monotonicity property. A simple lemma, the *Distinguishable Monotonicity Lema*, will then guarantee that

> *The set of undominated strategies of a player i with approximate-valuation set K_i consist of all valuations between the minimum integer and the maximum integer in K_i.*

[5]Actually relation (5.2) holds when the total number of coins usable by the players is bounded.

We believe that this simple property will be useful beyond our immediate need to prove Theorem 2.a and Theorem 3.a. Note that:

- Our setting is still discrete: continuous domains are only tools for proving the lemma.
- The Distinguishable Monotonicity Lemma, when specialized to the case where players know their valuations exactly, is a strengthening of a classical lemma that characterizes those mechanisms that are (very-weakly-)dominant-strategy-truthful in single-good auctions.
- The Distinguishable Monotonicity Lemma actually applies to *all* single-parameter domains, not just single-good auctions (the same way that the classical lemma does).

6 Our Model vs. Bayesian Ones in Single-Good Auctions

Observe that, in single-good auctions, if a player i knows that his true value θ_i is drawn from a distribution D_i, then, no matter what the auction mechanism M might be, he should essentially act as if his true valuation were *exactly* the expectation of D_i.[6] This implies that, for such auctions, our model subsumes or is equivalent to other Bayesian models.

Individual-Bayesian model as a special case. Assume that player i knows a distribution D_i from which is θ_i is drawn. By the above observation, in single-good auctions, this is equivalent to having an exact valuation, and of course the exact-valuation model is a special case of our model. Even if player i only knew the support D_i, in our model he could set this support to be the set K_i.

Equivalence to the Knightian model. Recall that in the Knightian model [Kni21, Bew02], a player i knows that his own true valuation is drawn from some distribution D_i belonging to a *set* of distributions $\mathscr{D}_i$, without being sure of *which* distribution in $\mathscr{D}_i$ is the right one. Thus, in single-good auctions, this model is equivalent to player i having the approximate-valuation set K_i, where $x \in K_i$ if and only if x is the expectation of some $D_i \in \mathscr{D}_i$. Automatically, therefore, any Knightian context is δ-approximate for some δ. Moreover, specific $\delta < 1$ can be explicitly computed from specific $\mathscr{D}_i$'s.

As an example, assume that each player i knows, for each possible valuation v, an interval $[A^i_v, B^i_v] \subseteq [0,1]$ where his probability of $\theta_i = v$ lies in, i.e., $\Pr[\theta_i = v] \in [A^i_v, B^i_v]$. If there exists δ such that for each player i and each valuation v, $[A^i_v, B^i_v] \subseteq \delta[x^i_v]$ for some x^i_v, then, although no player might know δ, after converting each $\mathscr{D}_i$ to K_i as described above, the context will be δ-approximate.

Equivalence to more general Bayesian models. For all contexts whose type spaces are "convex", and thus for single-good auction contexts, all "reasonable" models of uncertainty (including uncertainty about uncertainty, and so on) actually collapse to our proposed set-theoretic model.

7 Related Work

In settings of incomplete information, two types of uncertainty have been studied extensively: (1) the uncertainty of each player i about θ_{-i}, the type subprofile of his opponents, and (2) the uncertainty of a designer about the players' types. Notice, however, that *neither* type of uncertainty is the one we are interested in. As said, we focus *solely* on the uncertainty that each player i has about *his own* payoff type θ_i.

[6] In essence, this follows from the linearity of expectation and the linearity of the utility function.

Bayesian models of "self uncertainty". Various works model the players' self uncertainty via probability distributions. Let us mention a few examples.

In single-good auctions, Milgrom [Mil89] studies the revenue difference between second-price and English auctions, when the players do not exactly know their valuations, but only that they are drawn from a *common* distribution.

Sandholm [San00] presents an example of an auction (with an unconventionally-defined utility function) where a player's valuation is drawn from the uniform distribution over $[0, 1]$, and argues that reporting the expected valuation (i.e., 0.5) is no longer dominant-strategy.

Porter et al. [PRST08] consider a scheduling problem where tasks are to be assigned to players, and each player i privately knows that he would fail to perform task j with probability p_{ij}. This failure rate can be understood as a distribution of player's private type. Their paper studies efficient dominant-strategy mechanisms in this setting.

Feige and Tennenholtz [FT11] consider the problem of scheduling n players to use the same machine. Each player i has a task requiring time length l_i, but he does not know l_i: he only knows that l_i is drawn from a distribution L_i. The authors study dominant-strategy mechanisms without monetary transfer, and prove that even if L_i's support has two elements, then no constant fraction of the maximum social welfare can be guaranteed. To overcome this difficulty, they introduce a different measure of social efficiency, which they call "fair share", and provide mechanisms to guarantee an $\Omega(1)$ fair share.[7]

Set-theoretic models of "self uncertainty". As already mentioned, in his work (further formalized by Bewley [Bew02]) Knight [Kni21] considers a player that has a set of distributions and knows that his true type is drawn from one of them. The Knightian model immediately implies that the preferences for a player are no longer *completely ordered*: some pairs of preferences may become incomparable. (For single-good auctions, as already argued in Section 6, it is equivalent to each player i knowing a set of candidates K_i for his true valuation.)

Most of the works in the Knightian model address decision making. Some authors, such as Aumann [Aum62], Dubra et al. [DMO04], Ok [Ok02] and Nascimento [Nas11] work with incomplete orders. Others authors discuss various ways of bypassing the set-theoretic component of the Knightian model by computing a single number from the set of expected utilities: [Dan10] picks the average or an arbitrary one, [Sch89] picks the so-called Choquet expectation, and [GS89] picks the maximum.

General equilibrium models with unordered preferences have been considered by Mas-Colell [MC74], Gale and Mas-Colell [GMC75], Shafer and Sonnenschein [SS75], and Fon and Otani [FO79]. More recently, Rigotti and Shannon [RS05] characterize the set of equilibria in a financial market problem.[8]

Lopomo and the previous two authors [LRS09] also construct explicit mechanisms in a Knightian model, but for a single player. Specifically they consider the rent-extraction problem under two notions of implementation: 1) when reporting the truth is at least as good as any other strategy 2) when reporting the truth is not strictly eliminated in favor of another strategy. (Notice that, not envisaging other players, these are not notions of dominance, since the latter should take into account *all* strategy sub-profiles of other players.)

[7]A "ρ fair share" is a property such that each player i has at least ρ success rate if all other players share the same distribution as his l_i.

[8] A strategy profile is an equilibrium if no player can deviate and strictly benefit no matter which distribution is picked from his set. Notice that such an equilibrium is not a notion of dominance.

Appendix

A Single-Good Auctions in the Approximate-Valuation Model

As for any game, an auction can be thought as consisting of two parts, a context and a mechanism. (Because they do not affect our solution concepts, the players' beliefs are not part of our contexts.)

Contexts. For $\delta \in [0, 1)$, a *δ-approximate auction context* C consists of the following components.

- $N = \{1, 2, \ldots, n\}$, the set of *players.*
- $\{0, 1, \ldots, B\}$, the set of possible valuations for any player. B is the *valuation bound.*
- θ, the *true-valuation profile*, where each $\theta_i \in \{0, 1, \ldots, B\}$.
- $\Omega = (N \cup \{\bot\}) \times \mathbb{R}^N$, the set of *outcomes.* If $(a, P) \in \Omega$ is an outcome, then we refer to a as its *allocation* and to P as its profile of *prices.* (If $a \in N$ then player a wins the good; if $a = \bot$ then the good remains unallocated.)
- 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.
- K, the *δ-approximate-valuation profile*, where, for all i, $\theta_i \in K_i \subseteq \delta[x_i]$ for some x_i. Here $\delta[x] \stackrel{\text{def}}{=} [(1-\delta)x], (1+\delta)x] \cap \{0, 1, \ldots, B\}$.

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

Knowledge. Unless otherwise specified, in a context $C = (n, B, \delta, \theta, K)$ each player i only knows K_i and that $\theta_i \in K_i$, but has no other information. (A mechanism designer only knows n and B, and possibly also δ when specified.)

Notation.

- $\mathscr{C}^\delta_{n,B}$ is the set of all δ-approximate auction contexts with n players and valuation bound B.
- the *social welfare* (function) SW is defined as $\mathrm{SW}(\theta, (a, P)) \stackrel{\text{def}}{=} \theta_a$ for every true-valuation profile θ and outcome (a, P).
- the *maximum* social welfare of a true-valuation profile θ, $\mathrm{MSW}(\theta)$, is defined to be $\max_{i \in N} \theta_i$.

Mechanisms. While our contexts have K as a new component, our mechanisms are finite and ordinary. Indeed a mechanism for $\mathscr{C}^\delta_{n,B}$ is a pair $M = (S, F)$ where

- $S = S_1 \times \cdots \times S_n$, where each S_i, the set of i's pure *strategies* under M, is finite and non-empty; and
- $F\colon S \to (N \cup \{\bot\}) \times \mathbb{R}^N$ is M's (possibly probabilistic) *outcome function.*

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 $F^A_i(s)$ and $F^P_i(s)$ 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 $F^A_i(\sigma) \stackrel{\text{def}}{=} \mathbb{E}_{s\sim\sigma}\big[F^A_i(s)\big]$ and $F^P_i(\sigma) \stackrel{\text{def}}{=} \mathbb{E}_{s\sim\sigma}\big[F^P_i(s)\big]$.
- We refer to F^A as the *allocation function of M*. More generally, an *allocation function* is a function $f\colon S \to [0, 1]^N$ such that, for all strategy profile $s \in S$, $\sum_{i \in N} f_i(s) \leq 1$.

B Dominance in the Approximate-Valuation Model

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

Definition B.1. *In a mechanism $M = (S, F)$ for some class of contexts $\mathscr{C}^{\delta}_{n,B}$, fix a player $i \in N$ with approximate-valuation set K_i. For a (possibly mixed) strategy $\sigma_i \in \Delta(S_i)$ and a pure strategy $s_i \in S_i$, we say that*

- σ_i **very-weakly dominates** s_i, *in symbols* $\sigma_i \underset{i,K_i}{\overset{\text{vw}}{\succ}} s_i$, *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* $\sigma_i \underset{i,K_i}{\overset{\text{w}}{\succ}} s_i$, *if*

$$\sigma_i \underset{i,K_i}{\overset{\text{vw}}{\succ}} s_i \text{ and }\;\; \exists\, \theta_i \in K_i\,,\, \exists\, t_{-i} \in S_{-i}\; :\; \mathbb{E}u_i(\theta_i, F(\sigma_i, t_{-i})) > \mathbb{E}u_i(\theta_i, F(s_i, t_{-i}))\;\; .$$

For K_i, the set of **(very-weakly-)dominant** *and* **undominated strategies** *are*

$$\mathsf{Dnt}_i(K_i) \stackrel{\text{def}}{=} \left\{ s_i \in S_i : \forall\, t_i \in S_i\,,\; s_i \underset{i,K_i}{\overset{\text{vw}}{\succ}} t_i \right\} \text{ and } \mathsf{UDed}_i(K_i) \stackrel{\text{def}}{=} \left\{ s_i \in S_i : \nexists\, \sigma_i \in \Delta(S_i) \text{ s.t. } \sigma_i \underset{i,K_i}{\overset{\text{w}}{\succ}} s_i \right\}.$$

Finally we set $\mathsf{Dnt}(K) \stackrel{\text{def}}{=} \mathsf{Dnt}_1(K_1) \times \cdots \times \mathsf{Dnt}_n(K_n)$ *and* $\mathsf{UDed}(K) \stackrel{\text{def}}{=} \mathsf{UDed}_1(K_1) \times \cdots \times \mathsf{UDed}_n(K_n)$.

Remark B.2.

- The above extensions of the classical notions are quite straightforward. Some attention must be paid only to the extension of weak dominance in order to maintain the original "intent". 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 $\theta_i \in K_i$.
- Our impossibility results for implementation in dominant strategies are very strong because they already apply relative to *very*-weakly-dominant ones. Thus, for simplicity, we use the notation Dnt_i instead of, say, "VWDnt_i" (and have no need to define weakly-dominant or even strictly-dominant ones).
- Our results about implementation in undominated strategies are (as for the classical setting) relative to *weak* dominance. Thus, for simplicity and tradition sake, we use the notation UDed_i instead of, say, "UWDed_i".[9]

Notice that we obviously have that:

Fact B.3. $\mathsf{UDed}_i(K_i) \neq \emptyset$ *for all* K_i.

[9] We choose to analyze only strategies that are not weakly dominated because, when considering instead *very-weak* dominance, two "equivalent" strategies may eliminate each other and the set UDed may become empty.

C Proof of Theorem 1

Theorem 1. *For all n, $\delta \in (0,1)$, $B > \frac{3-\delta}{2\delta}$, and (possibly probabilistic) very-weakly-dominant-strategy-truthful mechanisms $M = (S, F)$, there exist a δ-approximate-valuation profile K and a true-valuation profile $\theta \in K$ 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, if a player reports a δ-integer-interval whose center is sufficiently high, then his winning probability and expected price remain constant.

Claim C.1. *For all players i, integers $x \in (\frac{3-\delta}{2\delta}, B]$, and δ-approximate-valuation sub-profiles $\widetilde{K}_{-i}$,*

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

and

$$F_i^P(\delta[x], \widetilde{K}_{-i}) = F_i^P(\delta[x+1], \widetilde{K}_{-i}) \ . \tag{C.2}$$

Proof of Claim C.1. Because the approximate-valuation set K_i of player i can be $\delta[x]$, and because when this is the case reporting the truth $\delta[x]$ very-weakly dominates $\delta[x+1]$, the following inequality must hold: $\forall\, \theta_i \in \delta[x]$,

$$F_i^A(\delta[x], \widetilde{K}_{-i}) \cdot \theta_i - F_i^P(\delta[x], \widetilde{K}_{-i}) \geq F_i^A(\delta[x+1], \widetilde{K}_{-i}) \cdot \theta_i - F_i^P(\delta[x+1], \widetilde{K}_{-i}) \tag{C.3}$$

Because K_i of player i can also be $\delta[x+1]$, and when this is the case reporting the truth $\delta[x+1]$ very-weakly dominates $\delta[x]$, the following inequality also holds: $\forall\, \theta'_i \in \delta[x+1]$,

$$F_i^A(\delta[x+1], \widetilde{K}_{-i}) \cdot \theta'_i - F_i^P(\delta[x+1], \widetilde{K}_{-i}) \geq F_i^A(\delta[x], \widetilde{K}_{-i}) \cdot \theta'_i - F_i^P(\delta[x], \widetilde{K}_{-i}) \ . \tag{C.4}$$

Thus, setting $\theta_i = x$ in Equation C.3 and $\theta'_i = x+1$ in Equation C.4, and summing up the resulting inequalities, the F_i^P price terms and a few other terms cancel out yielding the following inequality:

$$F_i^A(\delta[x+1], \widetilde{K}_{-i}) \geq F_i^A(\delta[x], \widetilde{K}_{-i}) \ . \tag{C.5}$$

Also, setting $\theta_i = \lfloor x(1+\delta)\rfloor$ in Equation C.3 and $\theta'_i = \lceil (x+1)(1-\delta)\rceil$ in Equation C.4,[10] and summing up the resulting inequalities we obtain the following one:

$$\Big(F_i^A(\delta[x], \widetilde{K}_{-i}) - F_i^A(\delta[x+1], \widetilde{K}_{-i})\Big) \cdot \Big(\lfloor x(1+\delta)\rfloor - \lceil (x+1)(1-\delta)\rceil\Big) \geq 0 \ . \tag{C.6}$$

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 Equation C.6 we deduce

$$F_i^A(\delta[x], \widetilde{K}_{-i}) \geq F_i^A(\delta[x+1], \widetilde{K}_{-i}) \tag{C.7}$$

Together, Equation C.5 and Equation C.7 imply the desired Equation C.1. Finally, combining Equation C.1 with Equation C.3 and Equation C.4 we obtain the desired Equation C.2. □

[10] 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 $\theta'_i \in \delta[x+1]$.

Let us now finish the proof of Theorem 1.

Choose the profile of approximate-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 C.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 approximate-valuation profile of the players is $K \stackrel{\text{def}}{=} (\delta[B], \delta[c], \ldots, \delta[c])$. Then, for the choice of true-valuation profile $\theta = (B, c, \ldots, c) \in K$, the expected social welfare is:

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

as desired. ∎

D The Undominated Intersection Lemma

Lemma D.1 (Undominated Intersection Lemma). *Let $M = (S, F)$ be a mechanism, i a player, and K_i and $\widetilde{K}_i$ two approximate-valuation sets of i intersecting in at least two integers. 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*

$$\forall\, s_{-i} \in S_{-i}\,, \quad \begin{aligned} &\left|F_i^A(\sigma_i, s_{-i}) - F_i^A(\widetilde{\sigma}_i, s_{-i})\right| < \varepsilon \\ &\left|F_i^P(\sigma_i, s_{-i}) - F_i^P(\widetilde{\sigma}_i, s_{-i})\right| < \varepsilon \end{aligned}$$

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 B.3, $\mathsf{UDed}_i(K_i)$ 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 $\widetilde{\sigma}_i \underset{i,\widetilde{K}_i}{\overset{\mathrm{w}}{\succ}} s_i$. We now argue that

$$\exists\, \tau_i \in \Delta(\mathsf{UDed}_i(K_i)) \text{ such that } \tau_i \underset{i,K_i}{\overset{\mathrm{w}}{\succ}} \widetilde{\sigma}_i \ .^{11} \tag{D.1}$$

Letting $\widetilde{\sigma}_i = \sum_{j \in X} \alpha^{(j)} s_i^{(j)}$ —where X is a subset of S_i— and 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 $\tau_i^{(j)} \underset{i,K_i}{\overset{\mathrm{w}}{\succ}} \widetilde{s}_i^{(j)}$. Thus, $\tau_i \stackrel{\text{def}}{=} \sum_{j \in X} \alpha^{(j)} \tau_i^{(j)}$ satisfies Equation D.1.

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

[11] 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 is the case in "$\tau_i \underset{i,K_i}{\overset{\mathrm{w}}{\succ}} \widetilde{\sigma}_i$" in Equation D.1.

infinite chain of (not necessarily distinct) strategies, $\{\sigma_i^{(k)}, \widetilde{\sigma}_i^{(k)}\}_{k\in\mathbb{N}}$, where:

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

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

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

Because $x_i \in K_i \cap \widetilde{K}_i$, setting $\theta_i = \widetilde{\theta}_i = x_i$ we see that, for all $s_{-i} \in S_{-i}$ and all $k \in \mathbb{N}$

$$\begin{aligned}
F_i^A(\sigma_i^{(k)}, s_{-i})x_i \quad &- \quad F_i^P(\sigma_i^{(k)}, s_{-i}) \qquad \leq \\
F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})x_i \quad &- \quad F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) \qquad \leq \\
F_i^A(\sigma_i^{(k+1)}, s_{-i})x_i \quad &- \quad F_i^P(\sigma_i^{(k+1)}, s_{-i}) \ .
\end{aligned}$$

Now notice that, $s_{-i} \in S_{-i}$, the infinite and non-decreasing sequence

$$F_i^A(\sigma_i^{(1)}, s_{-i})x_i - F_i^P(\sigma_i^{(1)}, s_{-i}) \leq F_i^A(\widetilde{\sigma}_i^{(2)}, s_{-i})x_i - F_i^P(\widetilde{\sigma}_i^{(2)}, s_{-i}) \leq F_i^A(\sigma_i^{(3)}, s_{-i})x_i - F_i^P(\sigma_i^{(3)}, s_{-i}) \leq \cdots$$

is upperbounded by B. (Indeed, $x_i \leq B$, F_i^A ranges between 0 and 1, and each price is non-negative.) Thus, by the Bolzano-Weierstrass theorem, for every $s_{-i} \in S_{-i}$ there must exist some $H_\varepsilon^{(s_{-i},x_i)} \in \mathbb{N}$ such that $\forall\, k > H_\varepsilon^{(s_{-i},x_i)}$:

$$\begin{aligned}
F_i^A(\sigma_i^{(k)}, s_{-i})x_i - F_i^P(\sigma_i^{(k)}, s_{-i}) \quad &\leq \quad F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) && \text{(D.2)} \\
F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) \quad &\leq \quad F_i^A(\sigma_i^{(k+1)}, s_{-i})x_i - F_i^P(\sigma_i^{(k+1)}, s_{-i}) && \text{(D.3)} \\
F_i^A(\sigma_i^{(k+1)}, s_{-i})x_i - F_i^P(\sigma_i^{(k+1)}, s_{-i}) \quad &\leq \quad F_i^A(\sigma_i^{(k)}, s_{-i})x_i - F_i^P(\sigma_i^{(k)}, s_{-i}) + \frac{\varepsilon}{3B} \ . && \text{(D.4)}
\end{aligned}$$

Similarly, because $y_i \in K_i \cap \widetilde{K}_i$, setting $\theta_i = \widetilde{\theta}_i = y_i$, we have that for every $s_{-i} \in S_{-i}$ there must exist some $H_\varepsilon^{(s_{-i},y_i)} \in \mathbb{N}$ such that $\forall\, k > H_\varepsilon^{(s_{-i},y_i)}$:

$$\begin{aligned}
F_i^A(\sigma_i^{(k)}, s_{-i})y_i - F_i^P(\sigma_i^{(k)}, s_{-i}) \quad &\leq \quad F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})y_i - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) && \text{(D.5)} \\
F_i^A(\widetilde{\sigma}_i^{(k)}, s_{-i})y_i - F_i^P(\widetilde{\sigma}_i^{(k)}, s_{-i}) \quad &\leq \quad F_i^A(\sigma_i^{(k+1)}, s_{-i})y_i - F_i^P(\sigma_i^{(k+1)}, s_{-i}) && \text{(D.6)} \\
F_i^A(\sigma_i^{(k+1)}, s_{-i})y_i - F_i^P(\sigma_i^{(k+1)}, s_{-i}) \quad &\leq \quad F_i^A(\sigma_i^{(k)}, s_{-i})y_i - F_i^P(\sigma_i^{(k)}, s_{-i}) + \frac{\varepsilon}{3B} \ . && \text{(D.7)}
\end{aligned}$$

Notice now that, because the set of strategies S_{-i} is finite, $H_\varepsilon = \max_{s_{-i}\in S_{-i}}\{H_\varepsilon^{(s_{-i},x_i)}, H_\varepsilon^{(s_{-i},y_i)}\}$ is a well defined integer. Next, we pick arbitrarily $\overline{k} > H_\varepsilon$, and prove that $\sigma_i^{(\overline{k}+1)}$ and $\widetilde{\sigma}_i^{(\overline{k})}$ are the two strategies that we are looking for.

To this end, pick arbitrarily $s_{-i} \in S_{-i}$ and sum up Equation D.2, Equation D.4 and Equation D.6. The (expected) prices and the $F_i^A(\sigma_i^{(\overline{k})}, s_{-i})x_i$ term will cancel out so as to yield

$$F_i^A(\sigma_i^{(\overline{k}+1)}, s_{-i})(x_i - y_i) \leq F_i^A(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i})(x_i - y_i) + \frac{\varepsilon}{3B} \ .$$

Then, sum up Equation D.3, Equation D.5 and Equation D.7. The (expected) prices and the $F_i^A(\sigma_i^{(\overline{k})}, s_{-i})y_i$ term will cancel out yielding

$$F_i^A(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i})(x_i - y_i) \leq F_i^A(\sigma_i^{(\overline{k}+1)}, s_{-i})(x_i - y_i) + \frac{\varepsilon}{3B} \ .$$

Since $x_i - y_i \geq 1$, we conclude that for all $s_{-i} \in S_{-i}$:

$$\left|F_i^A(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i}) - F_i^A(\sigma_i^{(\overline{k}+1)}, s_{-i})\right| \leq \frac{\varepsilon}{3B} \leq \varepsilon \ . \tag{D.8}$$

That is, the first inequality of Lemma D.1 has been proved. Let us now consider the price terms.

Fixing arbitrarily $s_{-i} \in S_{-i}$ and combining Equation D.3 and Equation D.8, we get:

$$\begin{aligned} F_i^A(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i})x_i - F_i^P(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i}) &\leq F_i^A(\sigma_i^{(\overline{k}+1)}, s_{-i})x_i - F_i^P(\sigma_i^{(\overline{k}+1)}, s_{-i}) \\ &\leq \Big(F_i^A(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i}) + \frac{\varepsilon}{3B}\Big)x_i - F_i^P(\sigma_i^{(\overline{k}+1)}, s_{-i}) \\ &\Rightarrow -F_i^P(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i}) \leq \frac{\varepsilon}{3} - F_i^P(\sigma_i^{(\overline{k}+1)}, s_{-i}) \ . \end{aligned} \tag{D.9}$$

Summing up Equation D.2 and Equation D.4 and then substituting Equation D.8, we get:

$$\begin{aligned} F_i^A(\sigma_i^{(\overline{k}+1)}, s_{-i})x_i - F_i^P(\sigma_i^{(\overline{k}+1)}, s_{-i}) &\leq F_i^A(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i})x_i - F_i^P(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i}) + \frac{\varepsilon}{3B} \\ &\leq (F_i^A(\sigma_i^{(\overline{k}+1)}, s_{-i}) + \frac{\varepsilon}{3B})x_i - F_i^P(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i}) + \frac{\varepsilon}{3B} \\ &\Rightarrow -F_i^P(\sigma_i^{(\overline{k}+1)}, s_{-i}) \leq \frac{2\varepsilon}{3} - F_i^P(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i}) \ . \end{aligned} \tag{D.10}$$

Finally, combining inequalities Equation D.9 and Equation D.10 we immediately get that for all $s_{-i} \in S_{-i}$:

$$\left|F_i^P(\widetilde{\sigma}_i^{(\overline{k})}, s_{-i}) - F_i^P(\sigma_i^{(\overline{k}+1)}, s_{-i})\right| \leq \frac{2\varepsilon}{3} \leq \varepsilon \ .$$

That is, the second desired inequality also holds, completing the proof of Lemma D.1. ■

E Proof of Theorem 2.b

Theorem 2.b. *Let $M = (S, F)$ be a* deterministic *mechanism. Then, for all $n, \delta \in (0, 1)$ and $B \geq \frac{1+\delta}{\delta}$, there exist a δ-approximate-valuation profile K, a true-valuation profile $\theta \in K$, and a strategy profile $s \in \mathsf{UDed}(K)$ such that:*

$$\mathrm{SW}\big(\theta, F(s)\big) \leq \left(\left(\frac{1-\delta}{1+\delta}\right)^2 + \frac{3}{B}\right)\mathrm{MSW}(\theta) \ . \tag{E.1}$$

Proof. Choose $x \stackrel{\text{def}}{=} \frac{B}{1+\delta}$ and $y \stackrel{\text{def}}{=} \frac{(1-\delta)x+2}{1+\delta}$, we have $x \geq y$ due to our choice of $B \geq \frac{1+\delta}{\delta}$. 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 $\lceil (1-\delta)x \rceil$ and $\lceil (1-\delta)x \rceil + 1$,[12] satisfying the requirement of (the Undominated Intersection) Lemma D.1.

Choose ε such that $\frac{1}{n} + \varepsilon < 1$. Then Lemma D.1 guarantees that

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

$$\left|F_i^A(\sigma_i, s_{-i}) - F_i^A(\sigma'_i, s_{-i})\right| < \varepsilon \ . \tag{E.2}$$

[12] We have $\lceil (1-\delta)x \rceil + 1 \leq x(1-\delta) + 2 = \frac{1-\delta}{1+\delta}B + 2 \leq B = (1+\delta)x$ as $B \geq \frac{1+\delta}{\delta}$, and therefore $\delta[x]$ contains both points. We also have $\lceil (1-\delta)x \rceil \geq \lceil (1-\delta)y \rceil$ and $\lceil (1-\delta)x \rceil + 1 \leq \lfloor (1-\delta)x \rfloor + 2 = \lfloor (1+\delta)y \rfloor$, and therefore $\delta[y]$ contains both points.

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 $F^A_j(\sigma'_1, \ldots, \sigma'_n) \le \frac{1}{n}$: 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 $F^A_1(\sigma'_1, s'_{-1}) \le \frac{1}{n}$. This together with Equation E.2 implies that $F^A_1(\sigma_1, s'_{-1}) \le \frac{1}{n} + \varepsilon < 1$. In turn, this implies that there exists a *pure* strategy $s_1 \in \mathsf{UDed}_1(\delta[x])$ such that, setting $s \stackrel{\text{def}}{=} (s_1, s'_{-1})$, $F^A_1(s) = 0$.

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

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

Note that $s \in \mathsf{UDed}(K)$, $\theta \in K$, and $\mathrm{MSW}(\theta) = (1+\delta)x = B$. Since $F^A_1(s) = 0$,

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

Thus the theorem holds. ■

F Proof of Theorem 3.b

Theorem 3.b. *Let $M = (S, F)$ be a (deterministic or probabilistic) mechanism. Then for all $n, \delta \in (0,1)$, and $B \ge \frac{1+\delta}{\delta}$, there exist a δ-approximate-valuation profile K, a strategy profile $s \in \mathsf{UDed}(K)$, and a true-valuation profile $\theta \in K$ such that*

$$\mathbb{E}\Big[\mathrm{SW}\big(\theta, F(s)\big)\Big] \le \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} + \frac{3}{B}\right) \mathrm{MSW}(\theta) \ . \tag{F.1}$$

Proof. (The first part of the proof closely tracks that of Theorem 2.a in Appendix E.[13])

Choose again $x \stackrel{\text{def}}{=} \frac{B}{1+\delta}$ and $y \stackrel{\text{def}}{=} \frac{(1-\delta)x+2}{1+\delta}$, we have $x \ge y$ due to our choice of $B \ge \frac{1+\delta}{\delta}$, and $\delta[x]$ and $\delta[y]$ both contain the two integers $\lceil (1-\delta)x \rceil$ and $\lceil (1-\delta)x \rceil + 1$, satisfying the requirement of (the Undominated Intersection) Lemma D.1.

Since we always have $\lceil (1-\delta)y \rceil < (1-\delta)y + 1$, we can choose $\varepsilon > 0$ small enough such that

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

Then (the Undominated Intersection) Lemma D.1 guarantees that

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

$$\big|F^A_i(\sigma_i, s_{-i}) - F^A_i(\sigma'_i, s_{-i})\big| < \varepsilon \ . \tag{F.2}$$

[13] 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 $F^A_1(s) = 0$.

Again consider the allocation distribution $F^A(\sigma'_1, \dots, \sigma'_n)$. By averaging, there exists some player j such that $F^A_j(\sigma'_1, \dots, \sigma'_n) \le \frac{1}{n}$. Without loss of generality let $j = 1$. Thus, by our choice of ε and Equation F.2, we have that $F^A_1(\sigma_1, \sigma'_{-1}) \le \frac{1}{n} + \varepsilon$. 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 $F^A_1(s_1, s'_{-1}) \le \frac{1}{n} + \varepsilon$. Now define

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

Notice that $s \in \mathsf{UDed}(K)$, $\theta \in K$, and $\mathrm{MSW}(\theta) = (1+\delta)x = B$. We now show that s, K, and θ satisfy the desired Equation F.1:

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

■

G The Distinguishable Monotonicity Lemma

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

Definition G.1. *If $f : [0, B]^N \to [0, 1]^N$ is an integrable*[14] *allocation function, then we denote by M_f the mechanism (S, F) where $S = \{0, 1, \dots, B\}^N$ 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 $j \ne i$.)*

Remark G.2.

- M_f is deterministic if and only if $f(\{0, 1, \dots, B\}) \subseteq \{0, 1\}^N$.
- For all player i and bid profile v, the expected price $F^P_i(v) = v_i \cdot f_i(v_i, v_{-i}) - \int_0^{v_i} f_i(z, v_{-i})\, dz$.
- We stress that M_f continues to have discrete strategy space $S = \{0, 1, \dots, B\}$, as the analysis over a continuous domain for f is only a tool for proving the lemma.
- Recall that an allocation function f is monotonic if each f_i is non-decreasing in the bid of player i, for any fixed choice of bids of all other players. In the exactly-valuation world, the class of mechanisms M_f's when f is both integrable and monotonic gives a full characterization to all (very-weakly-)dominant-strategy-truthful mechanisms in single-good auctions.

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

Now, we want to slightly strengthen this notion of monotonicity.

Definition G.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.

Distinguishable monotonicity is certainly an additional requirement to monotonicity, but actually quite mild. Indeed, the second-price mechanism is 2-DM and, if ties are broken at random, even 1-DM.[15] Yet, in our approximate-valuation world, this mild additional requirement is quite useful for "controlling" the *undominated strategies* of a mechanism, and thus for engineering implementations of desirable social choice functions in undominated strategies.

Lemma G.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-valuation profile K,*

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

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

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

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

[15] For example, the allocation function of the second-price mechanism with lexicographic tie-breaking is M_f, where $\forall i \in N$ and $\forall v \in \{0, \ldots, B\}^N$:

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

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

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. □

H Proof of Theorem 2.a

Theorem 2.a. *Let $M_{\mathsf{2P}} = (S_{\mathsf{2P}}, F_{\mathsf{2P}})$ be the second-price mechanism with a deterministic tie-breaking rule. Then, for all n, $\delta \in [0, 1)$, B, δ-approximate-valuation profiles K, true-valuation profiles $\theta \in K$, and strategy profiles $v \in \mathsf{UDed}(K)$:*

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

Proof. Let $M_{\mathsf{2P}} = (S_{\mathsf{2P}}, F_{\mathsf{2P}})$ be (a version of) the second-price mechanism with a *deterministic* tie-breaking rule. Since K is a δ-approximate-valuation set, for each player i let x_i be such that $K_i \subseteq \delta[x_i] \cap \{0, \ldots, B\}$. Then, in light of (the Distinguishably Monotonicity) Lemma G.4 and the previous observation that F_{2P}^A is a 2-DM allocation function, we have that, for each player i:

$$\mathsf{UDed}_i(x) \subseteq \big\{\lceil(1-\delta)x_i\rceil - 1, \ldots, \lfloor(1+\delta)x_i\rfloor + 1\big\} \ . \tag{H.2}$$

Now we prove the lower bound to the social welfare. Let $\theta \in K$ be a candidate true-valuation profile, i^* the player with the highest valuation according to θ, and j^* the player winning the good under the bid profile v. (Thus, $\theta_{i^*} = \max_i \theta_i$ and $v_{j^*} = \max_j v_j$.) We now bound the difference between θ_{i^*} and θ_{j^*} when $i^* \neq j^*$.

From Equation H.2 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 $(1-\delta)x_{i^*} \leq (1+\delta)x_{j^*} + 2$. Since we also know that $\theta_{j^*} \geq (1-\delta)x_{j^*}$ and $(1+\delta)x_{i^*} \geq \theta_{i^*}$, we obtain:

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

Thus, the claim of our theorem holds. ■

Consider the case where the second-price mechanism breaks ties at random (assigning a positive probability to each tie). Then, one can use a proof analogous to the one above, with the only difference being that F_{2P}^A is 1-DM (instead of only 2-DM), and invoking the stronger inclusion of (the Distinguishably Monotonicity) Lemma G.4, to show the following, stronger lower bound:

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

I Proof of Theorem 3.a

In this section we explicitly construct and analyze the desired mechanism $M^{(\delta)}_{\mathsf{opt}}$. This process is not going to be trivial, and thus we break it into several steps, providing intuition as needed.

I.1 A Very Restricted Search

In order to leverage our Distinguishable Monotonicity Lemma (Lemma G.4), it is natural for us to search for $M^{(\delta)}_{\mathsf{opt}}$ among 1-DM mechanisms. Let us now distill an additional requirement for the underlying allocation function of such mechanisms that suffices for our goals. We shall do so in terms of the following positive quantity D_δ: for all $\delta \in (0,1)$,

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

Definition I.1. *We say that a allocation function f is δ-**good** if it is 1-DM and:*

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

The reason why the additional requirement is sufficient is easily understood:

Lemma I.2. *If f is δ-good, then M_f satisfies that such that for every δ-approximate-valuation profile K, every strategy profile $s \in \mathsf{UDed}(K)$ and every true-valuation profile $\theta \in K$:*

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

Proof. Let K be an arbitrarily chosen δ-approximate-valuation profile. Then, because in any allocation the social welfare coincides with the welfare of a given player, to prove our lemma it suffices to prove that

$$\forall \theta \in K, \quad \forall v \in \mathsf{UDed}(K), \quad \forall i \in N, \quad \sum_{j=1}^{n} \theta_j f_j(v) \geq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}\right)\theta_i\ . \tag{I.2}$$

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, $\theta \in K$ and the Distinguishable Monotonicity Lemma respectively imply

$$(1-\delta)x_i \leq \theta_i \leq (1+\delta)x_i \quad \text{and} \quad (1-\delta)x_i \leq \min K_i \leq v_i \leq \max K_i \leq (1+\delta)x_i.$$

Combining these two chains of inequalities yields

$$\frac{1-\delta}{1+\delta}v_i \leq \theta_i \leq \frac{1+\delta}{1-\delta}v_i\ . \tag{I.3}$$

Let us now argue that Equation I.2 holds by arbitrarily fixing v and i and showing that it is impossible to construct a "bad" θ so as to violate Equation I.2.

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 $f_i(v) \geq \big(\frac{(1-\delta)^2+\frac{4\delta}{n}}{(1+\delta)^2}\big)$, or as small as possible otherwise. So there are two extreme θ's.

Considering these extreme choices, we conclude that no θ contradicts Equation I.2 if:

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

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

Simplifying the above equations, Equation I.2 holds if both the following inequalities hold:

$$\sum_{j=1}^{n} v_j f_j(v) \geq \frac{n+D_\delta}{n}\cdot\frac{1}{D_\delta+1}\cdot v_i \ , \tag{I.4}$$

$$\sum_{j=1}^{n} v_j f_j(v) + D_\delta\cdot v_i f_i(v) \geq \frac{n+D_\delta}{n}v_i \ . \tag{I.5}$$

Note that Equation I.5 holds because it is implied by the hypothesis that f is δ-good; note also that Equation I.4 holds because it is implied by Equation I.5. 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 Equation I.2 and our lemma hold. □

I.2 Our Allocation Function

In light of our last lemma, all is left is to find a suitable δ-good allocation function f.

Some intuition. If the players' bids are not "clustered", then f should clearly give a much higher probability mass to the highest bids, as lower bids are less likely to come from players with high true valuations. However, when the highest bids are close to each other, it is hard for f to "infer" from them who the player with the highest true valuation really is — after all, we are in an approximate-valuation model. The intelligent thing for f to do in such a case is to assign the good to a randomly chosen high-bidding player. To achieve optimality, however, one must be much more careful in allocating probability mass, and some complexities should be expected.

Since the mechanism $M^{(\delta)}_{\mathsf{opt}}$ of Theorem 3.a is allowed to depend on the approximation accuracy δ, we construct its allocation function, $f^{(\delta)}$, depending on it. Our proposed $f^{(\delta)}$ derives from the players' bids a threshold, and probabilistically chooses the winning player only among those bids lying above the threshold. We now explain the rationale for these choices.

Recall that, to be δ-good, a allocation function $f:[0,B]^N\to[0,1]^N$ should satisfy Equation I.1, that is:

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

A reasonable guess to "solve for f" is to restrict our attention to *symmetric* functions. The most natural candidate is simply

$$\forall z\in[0,B]^N,\quad f_i(z)=\frac{1}{n}\cdot\frac{z_i(n+D_\delta)-\sum_{j=1}^{n} z_j}{z_i D_\delta} \ .$$

One could verify that the function f, in addition to being symmetric, sums up to 1, is 1-DM, and satisfies the desired condition Equation I.1. (In fact, as we shall see, the above candidate f coincides with our proposed $f^{(\delta)}$ when no threshold is introduced.) We would be done, except for one crucial fact: f sometimes takes negative values!

We therefore need to "patch" the guessed function f by forcing non-negativity while maintaining the other required properties, and this is exactly where the idea of a threshold, winners, and losers comes in. Roughly, only players with sufficiently low reported valuations are at risk of a "negative probability" and, because are most likely to have low true valuations, we remove them from the auction altogether. To preserve the other properties, though, we need to re-weight the function, thereby obtaining Equation I.6. Thus, at high level, we simply keep removing players until all of the players are given non-negative probability (by virtue of being in the auction or having been thrown out). A similar idea previously appeared in [CLS$^+$11].

While the introduction of a threshold fixes the "negativity problem", it introduces additional complexities. (For example, even the simple task of verifying monotonicity, where the bids of all players but i are fixed, becomes non-trivial. Indeed, the number of winners n^* varies as the bid of player i increases, and thus the definition of $f^{(\delta)}$ varies too.)

Let us now proceed more formally. Recall that $D_\delta \stackrel{\text{def}}{=} \left(\frac{1+\delta}{1-\delta}\right)^2 - 1 > 0$.

Definition I.3. *For every $\delta \in (0,1)$, define the function $f^{(\delta)} \colon [0,B]^N \to [0,1]^N$ 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*

$$f_i^{(\delta)}(z) \stackrel{\text{def}}{=} \begin{cases} \frac{1}{n} \cdot \frac{n+D_\delta}{n^*+D_\delta} \cdot \frac{z_i(n^*+D_\delta) - \sum_{j=1}^{n^*} z_j}{z_i D_\delta}, & \text{if } i \leq n^*, \\ 0, & \text{if } i > n^*; \end{cases} \quad , \tag{I.6}$$

  where $n^* \in \{1, 2, \ldots, n\}$ is the index in N (whose existence and uniqueness will be proved shortly) such that

$$z_1 \geq \cdots \geq z_{n^*} > \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta} \geq z_{n^*+1} \geq \cdots \geq z_n \quad . \tag{I.7}$$

- *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 $z_{\pi(1)} \geq \cdots \geq z_{\pi(n)}$ (i.e., we define $f_i^{(\delta)}$ by extending it symmetrically).*

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

Lemma I.4. *$f^{(\delta)}$ is a well-defined allocation function.*

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

Assume, without loss of generality, that $z_1 \geq z_2 \geq \cdots \geq z_n$. Note that there exists an index n' in N such that

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

Indeed, Equation I.8 vacuously holds for $n' = n$. Now take n'' to be the least such index. Accordingly,

$$\forall\, i > n'', \quad z_i \leq \frac{\sum_{j=1}^{n''} z_j}{n'' + D_\delta} \quad . \tag{I.8}$$

Next we claim that

$$\forall\, i \le n'', \quad z_i > \frac{\sum_{j=1}^{n''} z_j}{n'' + D_\delta} \ . \tag{I.9}$$

To prove Equation I.9, it suffices to consider $i = n''$ because z is non-increasing. Indeed, by the minimality of n'' we know that ("$n''-1$ does not work", that is) there exists some $j \ge n''$ such that

$$z_{n''} \ge z_j > \frac{\sum_{j=1}^{n''-1} z_j}{n'' - 1 + D_\delta} \ ,$$

which, after rearranging, is equivalent to $z_{n''} > \frac{\sum_{j=1}^{n''} z_j}{n''+D_\delta}$ as desired.

At last, combining Equation I.8 and Equation I.9, and choosing $n^* = n''$, Equation I.7 is satisfied.

Next, we prove that n^* is unique. Suppose by way of contradiction that there exist two integers $n^\perp$ and $n^\top$, with $n^\perp < n^\top$ both satisfying Equation I.7. Now define

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

By invoking Equation I.7 with $n^\top$ and $n^\perp$, we deduce that for $i \in \{n^\perp + 1, \ldots, n^\top\}$,

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

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

$$\frac{S^\perp}{n^\perp + D_\delta} \ge \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + S^\Delta}{n^\perp + n^\Delta + D_\delta} \ . \tag{I.10}$$

Let us now show that the second inequality of Equation I.10 contradicts the first inequality Equation I.10:

$$\begin{aligned} \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + S^\Delta}{n^\perp + n^\Delta + D_\delta} &\Leftrightarrow (n^\perp + n^\Delta + D_\delta) S^\Delta > n^\Delta (S^\perp + S^\Delta) \\ &\Leftrightarrow (n^\perp + D_\delta) S^\Delta > n^\Delta S^\perp \Leftrightarrow \frac{S^\Delta}{n^\Delta} > \frac{S^\perp}{(n^\perp + D_\delta)} \ . \end{aligned} \tag{I.11}$$

The contradiction establishes the uniqueness of n^*.

We are left to prove that (a) $f_i^{(\delta)}(z) \ge 0$ for every i and z, and (b) $\sum_i f_i^{(\delta)}(z) \le 1$ for every z. (Indeed, the last two properties imply that $f_i^{(\delta)}(z) \le 1$.)

Assume, again without loss of generality, that $z_1 \ge z_2 \ge \cdots \ge z_n$. Equation I.7 tells us that $z_i(n^* + D_\delta) - \sum_{j=1}^{n^*} z_j \ge 0$ for each $i \le n^*$, so (a) follows immediately. As for (b),

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

□

Lemma I.5. *$f^{(\delta)}$ is monotonic.*

Proof. By symmetry it suffices to show that $f^{(\delta)}$ is monotonic with respect to the n-th coordinate. Without loss of generality, assume $z_1 \geq z_2 \geq \cdots \geq z_{n-1}$. We need to prove that for any $z_n^{\perp}$ and $z_n^{\top}$ with $0 \leq z_n^{\perp} < z_n^{\top} \leq B$,

$$f_n^{(\delta)}(z_{-n}, z_n^{\perp}) \leq f_n^{(\delta)}(z_{-n}, z_n^{\top}) \ . \tag{I.12}$$

We will prove Equation I.12 in three steps.

- **Step 1.** Letting n' be the number of winners in a game where only the first $n-1$ players are bidding z_{-n}, we first prove that:

$$z_n \leq \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \longrightarrow f_n^{(\delta)}(z_{-n}, z_n) = 0 \ \text{(i.e., } n \text{ is a loser)} \tag{I.13}$$

$$z_n > \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \longrightarrow f_n^{(\delta)}(z_{-n}, z_n) > 0 \ \text{(i.e., } n \text{ is a winner)} \tag{I.14}$$

To show Equation I.13, recall that, in the game with only the first $n-1$ players bidding z_{-n}, we have n' winners satisfying,

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

Then imagine that player n comes with bid z_n that is at most $\frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$. In this new game, because the threshold does not change, the set of winners continues to be $\{1, 2, \ldots, n'\}$ and therefore n must be a loser. Indeed,

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

To show Equation I.14, we actually prove its contrapositive: namely,

$$f_n^{(\delta)}(z_{-n}, z_n) = 0 \text{ (i.e. } n \text{ is a loser)} \longrightarrow z \leq \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \ .$$

Let n^* be the number of winners when $f_n^{(\delta)}(z_{-n}, z_n) = 0$, that is, in the game where there are n players, the bid profile is z, and player n is a loser; then,

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

The above also implies the following, where player n has been removed:

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

This means, n^* is also the number of winners for the $(n-1)$-player game, i.e., $n^* = n'$. This gives $z_n \leq \frac{\sum_{j=1}^{n^*} z_j}{n^*+D_\delta} = \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$. □

Because of Step 1, we only need to show Equation I.12 for $z_n^\perp$ and $z_n^\top$ satisfying $z_n^\top > z_n^\perp > \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$. Notice that in such a case, player n is always a winner. Therefore, let $\{1,\dots,n^\perp,n\}$ and $\{1,\dots,n^\top,n\}$ be the winners when the bid profiles are $(z_{-n},z_n^\perp)$ and $(z_{-n},z_n^\top)$ respectively.

- **Step 2.** We now prove that

$$n^\perp \geq n^\top \ . \tag{I.15}$$

Assume by way of contradiction that $n^\perp < n^\top$ and. As in Lemma I.4, set $n^\Delta \stackrel{\text{def}}{=} n^\top - n^\perp$, $S^\perp \stackrel{\text{def}}{=} \sum_{j=1}^{n^\perp} z_j$ and $S^\top \stackrel{\text{def}}{=} \sum_{j=1}^{n^\top} z_j = S^\perp + S^\Delta$. Then each player i, with $n^\perp \leq i < n^\top$, is a loser when the bid profile is $(z_{-n},z_n^\perp)$ while a winner when the bid profile is $(z_{-n},z_n^\top)$; in particular,

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

Averaging over all $n^\perp \leq i < n^\top$ we get:

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

but this is already a contradiction, since the right hand side is equivalent to (using a similar technique as Equation I.11):

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\perp + S^\Delta + z_n^\top}{n^\perp + n^\Delta + 1 + D_\delta} \Leftrightarrow \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + z_n^\top}{n^\perp + 1 + D_\delta} \ ,$$

which actually contradicts the left hand side, as $z_n^\top > z_n^\perp$. Therefore, $n^\perp \geq n^\top$. □

We now use the fact that $n^\perp \geq n^\top$ to obtain Equation I.12 for such $z_n^\perp$ and $z_n^\top$ satisfying $z_n^\top > z_n^\perp > \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$.

- **Step 3.** We now prove Equation I.12.

If $n^\perp = n^\top$, then for both $(z_{-n},z_n^\top)$ and $(z_{-n},z_n^\perp)$, the set of winners is $\{1,2,\dots,n^\perp,n\}$. Let $n^* = n^\perp + 1 = n^\top + 1$ be the number of winners and 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 $n^\perp > n^\top$, 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. Then we average over all z_i for $n^\top < i \leq n^\perp$ and get:

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} = \frac{S^\top + S^\Delta + z_n^\perp}{n^\top + n^\Delta + 1 + D_\delta} \ . \tag{I.16}$$

But this is equivalent to (again using the same technique as Equation I.11)

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\top + z_n^\perp}{n^\top + 1 + D_\delta} \ . \tag{I.17}$$

Letting $C_1 = \frac{n+D_\delta}{n}$, we now do the final calculation:

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

Here the last inequality has used $z_n^\top - z_n^\perp \geq 0$ and $S^\Delta(n^\top + 1 + D_\delta) - n^\Delta(S^\top + z_n^\perp) > 0$ (by Equation I.17). □

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

Lemma I.6. *$f^{(\delta)}$ is 1-distinguishably monotonic.*

Proof. We already know from Lemma I.5 that $f^{(\delta)}$ is monotonic. Also, the integrability of $f^{(\delta)}$ is obvious, because $f^{(\delta)}$ is piecewise continuous, and there are at most n pieces, as the number of winners decreases when z_n increases (recall Equation I.15). 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 $v_i < v'_i$. Define $v_{-i} \stackrel{\text{def}}{=} (v_i, v_i, \ldots, v_i)$, then:

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

Here the upper bound $z \leq (1 + D_\delta)v_i$ is to make sure that the number of winners is still n on input (z, v_{-i}). Notice that $f(z, v_{-i})$ is a function that is *strictly* increasing when z increases in such range, and therefore

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

as desired. □

Lemma I.7. *$f^{(\delta)}$ is δ-good.*

Proof. We already know from Lemma I.6 that $f^{(\delta)}$ is 1-DM. Therefore, in order to prove that $f^{(\delta)}$ is δ-good, we only need to show that Equation I.1 holds. We will actually prove that Equation I.1 holds not only for the discrete cube $\{0, 1, \ldots, B\}^N$ but also in the continuous cube $[0, B]^N$.

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

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

For each player k with $k > n^*$, because he is a loser, we have,

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

satisfying Equation I.1, where the last inequality is due to $k > n^*$ and Equation I.7.

For each winner i (i.e., with $i \leq n^*$), we have

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

again satisfying Equation I.1. □

I.3 Our Mechanism $M_{\mathsf{opt}}^{(\delta)}$

Theorem 3.b. *$\forall n$, $\forall \delta \in (0,1)$, and $\forall B$, there exists a mechanism $M_{\mathsf{opt}}^{(\delta)}$ such that for every δ-approximate-valuation profile K, every true-valuation profile $\theta \in K$, and every strategy profile $v \in \mathsf{UDed}(K)$:*

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

Proof. By Lemma I.7, the function $f^{(\delta)}$ from Definition I.3 is a (well-defined) allocation function that is also δ-good. Therefore, invoking Lemma I.2, the mechanism $M_{\mathsf{opt}}^{(\delta)} \stackrel{\text{def}}{=} M_{f^{(\delta)}}$ yields the target social welfare. ■

Finally, we note that $M_{\mathsf{opt}}^{(\delta)}$ can be implemented efficiently (just like the second-price mechanism):

Claim I.8. *The outcome function F of $M_{\mathsf{opt}}^{(\delta)}$ is efficiently computable.*

Proof. It suffices to show that both the allocation function $F^A = f^{(\delta)}\big|_{\{0,1,\dots,B\}^N}$ and expected price function F^P are efficiently computable over $\{0,1,\dots,B\}^N$.

First, $f^{(\delta)}$ is efficiently computable for trivial reasons: the number of winners n^* is between 1 and n and can be determined in linear time.

However, F^P is efficiently computable for a more involved reason. Without loss of generality, we show how to compute the expected price for player n as a function of v_n, i.e.,

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

Indeed, when v_{-n} is fixed, $f_n^{(\delta)}$ is a function piece-wisely defined according with respect 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 line 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 intergrable. Therefore, the only question is how to calculate the pieces for $f_n^{(\delta)}$.

This is again not hard, by using a simple line sweep method. One can start from $v_n = \frac{\sum_{j=1}^{n'} v_j}{n'+D_\delta}$ and move v_n upwards. At any moment, one can calculate the earliest time that Equation I.7 is violated, and claim that another piece of $f_n^{(\delta)}$ is found. $\square$

J Performance Diagrams

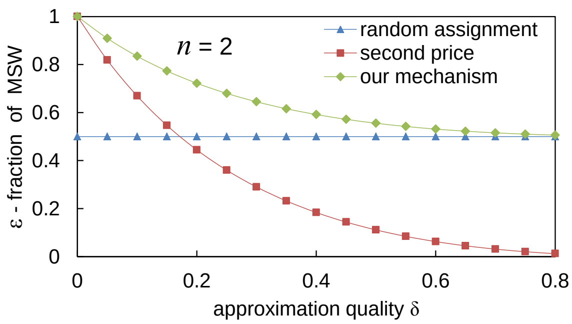

(a) With $n = 2$ players, the second-price mechanism performs *worse* than randomly assigning the good for $\delta > 0.18$.

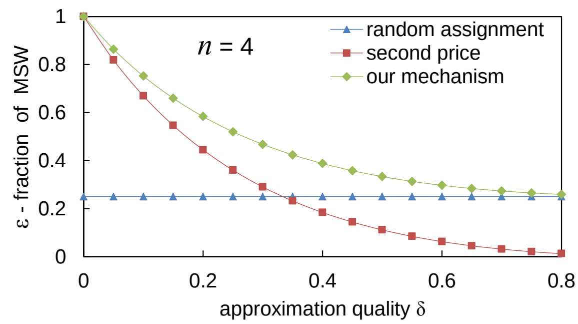

(b) With $n = 4$ players, the second-price mechanism performs *worse* than randomly assigning the good for $\delta > 0.34$.

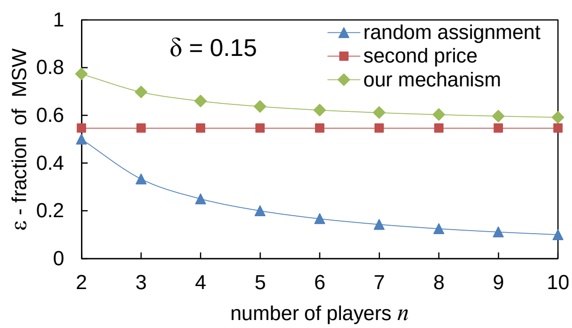

(c) With $\delta = 0.15$, the second-price mechanism always performs better than randomly assigning the good.

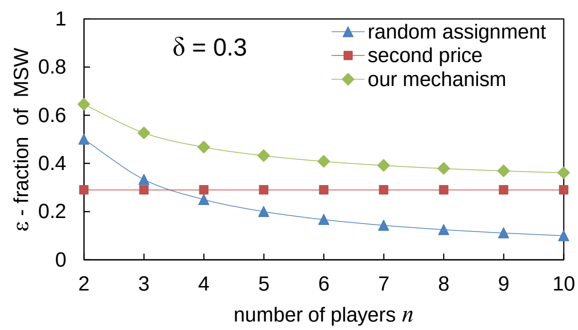

(d) With $\delta = 0.3$, the second-price mechanism performs *worse* than randomly assigning the good for $n = 2, 3$.

Figure 1: We compare the social welfare guarantees of randomly assigning the good ($\varepsilon = \frac{1}{n}$), the second-price mechanism ($\varepsilon = \frac{(1-\delta)^2}{(1+\delta)^2}$, see Theorem 2), and our optimal mechanism ($\varepsilon = \frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}$, see Theorem 3). 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

[Aum62] Robert J. Aumann. Utility theory without the completeness axiom. *Econometrica*, 30(3):445–462, July 1962. 9

[Bew02] 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. 1, 3, 8, 9

[CLS+11] Wei Chen, Pinyan Lu, Xiaorui Sun, Bo Tang, Yajun Wang, and Zeyuan Allen Zhu. Optimal pricing in social networks with incomplete information. In *WINE '11: The 7th Workshop on Internet & Network Economics*. Springer, 2011. 22

[Dan10] Eric Danan. Randomization vs. selection: How to choose in the absence of preference? *Management Science*, 56:503–518, March 2010. 9

[DMO04] 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. 9

[FO79] 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. 9

[FT11] 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. arXiv:1103.2520v1. 1, 9

[GMC75] 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. 9

[GS89] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, April 1989. 9

[Jac92] Matthew O Jackson. Implementation in undominated.strategies: A look at bounded mechanisms. *Review of Economic Studies*, 59(4):757–75, October 1992. 3

[Kni21] Frank H. Knight. *Risk, Uncertainty and Profit*. Houghton Mifflin, 1921. 1, 3, 8, 9

[LRS09] Giuseppe Lopomo, Luca Rigatti, and Chris Shannon. Uncertainty in mechanism design, 2009. `http://www.pitt.edu/~luca/Papers/mechanismdesign.pdf`. 9

[MC74] Andrew Mas-Colell. An equilibrium existence theorem without complete or transitive preferences. *Journal of Mathematical Economics*, 1(3):237–246, December 1974. 9

[Mil89] Paul Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22, Summer 1989. 9

[Mye81] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. 1, 4

[Nas11] Leandro Nascimento. Remarks on the consumer problem under incomplete preferences. *Theory and Decision*, 70(1):95–110, January 2011. 9

[Ok02] Efe A. Ok. Utility representation of an incomplete preference relation. *Journal of Economic Theory*, 104:429–449, 2002. 9

[PRST08] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172:1783–1799, October 2008. 9

[RS05] Luca Rigotti and Chris Shannon. Uncertainty and risk in financial markets. *Econometrica*, 73(1):203–243, 01 2005. 9

[San00] Tuomas Sandholm. Issues in computational vickrey auctions. *International Journal of Electronic Commerce*, 4:107–129, March 2000. 1, 9

[Sch89] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–87, May 1989. 9

[SS75] Wayne Shafer and Hugo Sonnenschein. Equilibrium in abstract economies without ordered preferences. *Journal of Mathematical Economics*, 2(3):345–348, December 1975. 9