Knightian Self Uncertainty in the VCG Mechanism for Unrestricted Combinatorial Auctions *

Alessandro Chiesa† alexch@csail.mit.edu MIT CSAIL

Silvio Micali† silvio@csail.mit.edu MIT CSAIL

Zeyuan Allen Zhu† zeyuan@csail.mit.edu MIT CSAIL

June 11, 2014

Abstract

We study the social welfare performance of the VCG mechanism in the well-known and challenging model of self uncertainty initially put forward by Frank H. Knight and later formalized by Truman F. Bewley. Namely, the only information that each player i has about his own true valuation consists of a set of distributions, from one of which i's valuation has been drawn.

We assume that each player knows his true valuation up to an additive inaccuracy δ, and study the social welfare performance of the VCG mechanism relative to $\delta > 0$. In this paper, we focus on the social welfare performance of the VCG mechanism in *unrestricted combinatorial auctions*, both in undominated strategies and regret-minimizing strategies. Denote by MSW the maximum social welfare.

Our first theorem proves that, in an n-player m-good combinatorial auction, the VCG mechanism may produce outcomes whose social welfare is $\leq \mathrm{MSW} - \Omega(2^m \delta)$, even when $n = 2$ and each player chooses an *undominated strategy*. We also geometrically characterize the set of undominated strategies in this setting.

Our second theorem shows that the VCG mechanism performs well in *regret-minimizing strategies*: the guaranteed social welfare is $\geq \mathrm{MSW} - 2\min\{m, n\}\delta$ if each player chooses a pure regret-minimizing strategy, and $\geq \mathrm{MSW} - O(n^2\delta)$ if mixed strategies are allowed.

Finally, we prove a lemma bridging two standard models of rationality: utility maximization and regret minimization. A special case of our lemma implies that, in any game (Knightian or not), every implementation for regret-minimizing players also applies to utility-maximizing players who use regret *only* to break ties among their undominated strategies. This bridging lemma thus implies that the VCG mechanism continues to perform very well also for the latter players.

---

*The one-paged abstract of this paper has appeared in the 15th ACM Conference on Economics and Computation (ACM-EC 2014). The weaker versions of our first, second theorems and the bridging lemma have respectively been organized as three independent technical reports [CMZ14b], [CMZ14c], and [CMZ14a].

†The authors would like to thank Jelani Nelson, Rafael Pass and Luca Trevisan for helpful discussions, the Skolkovo Foundation (under agreement dated 10/26/2011), and two Simons Graduate Student Awards (under grant no. 282907, 284059).

Contents

1 Introduction

In a private-value auction, a *valuation* is a function mapping each possible allocation of the good(s) to a real number. Private-value auctions are traditionally studied in the *exact-valuation* model, that is, assuming that every player i knows his own true valuation, θ_i^*, *exactly.* This assumption cannot hold in all concrete applications. We thus wish to investigate what happens to a classical auction mechanism, the VCG, when some (or every) player i is uncertain about θ_i^*.

A traditional way to capture such self uncertainty is the 'single-distribution' model, which assumes that a player i only knows the true distribution D_i (over the set of all his possible valuations, Θ_i) from which θ_i^* has been drawn. This assumption, however, may be too strong. Sometimes, a player i may not even be able to establish which of two candidate valuations is 'more likely' to be θ_i^*. We thus study the VCG mechanism under a more general model, envisaged by Frank H. Knight almost a century ago, and later formalized by Truman F. Bewley.

The Knightian Auction Model. Informally, focusing on auctions and letting θ^* be the profile of true valuations, for each player i,

i's sole information about θ^* consists of $\mathcal{K}_i$, **a set of distributions** *over Θ_i, from one of which θ_i^* has been drawn. (The true valuations are uncorrelated.)*

That is, $\mathcal{K}_i$ is the sole (and private) information that i has about his own true valuation θ_i^*. Furthermore, for every opponent j, i has no information (or beliefs) about θ_j^* or $\mathcal{K}_j$. Therefore, not only is our model one of extreme *incomplete information*, but also one of *incomplete preferences.*

Let us give two simple examples of Knightian auctions of a single good.

Example 1. A player i knows that θ_i^* is drawn from a Gaussian distribution $N(\mu, \sigma)$ over $\mathbb{R}$, but not the mean, μ, nor the standard deviation, σ, of $N(\mu, \sigma)$. He only knows that $\mu = 10 \pm 1$ and $\sigma = 2 \pm 1$. Thus, $\mathcal{K}_i = \{N(\mu, \sigma) : \mu \in [9, 11], \sigma \in [1, 3]\}$.

Example 2. The good is an exclusive licence to a cryptographic algorithm whose security is based on unproven mathematical assumption A. Then a player i's valuation may be drawn from a distribution D_0 if A is false, and from a distribution D_1 if A is correct. Thus, $\mathcal{K}_i = \{D_0, D_1\}$.

A Mathematically Equivalent Model. For the auctions we consider (single-good, multi-unit, and unrestricted-combinatorial), an *equivalent* formulation of the Knightian model is the following: for each player i,

i's sole information about θ^* consists of $K_i \subseteq \Theta_i$,
a set of valuations *that includes θ_i^*.*

We refer to K_i as the *candidate (valuation) set* of player i.

The equivalence between the latter and the previous formulation is due to the fact that a player i may 'collapse' each distribution $D_i \in \mathcal{K}_i$ to its expectation $\mathbb{E}(D_i)$, given that all he cares about is his expected (quasi-linear) utility.[1]

Example: Consider an auction of two goods, a and b. In this case, a player's valuation consists of a triple $(v_a, v_b, v_{\{a,b\}})$, where v_a is his value for a alone, v_b for b alone, and $v_{\{a,b\}}$ for a and b together. Suppose a player i knows that θ_i^* is drawn from a distribution D_i over such triples, and let $(e_a, e_b, e_{\{a,b\}}) \stackrel{\text{def}}{=} \mathbb{E}_{\theta_i \sim D_i}\big[\big(\theta_i(a), \theta_i(b), \theta_i(\{a,b\})\big)\big]$. Now, consider any outcome ω of the auction. Without loss of generality, let ω be such that i wins good a and pays price

[1] Indeed, whatever the auction mechanism used, this equivalence holds for any auction where each Θ_i is a *convex* set.

p. Then, i's expected utility for ω is $\mathbb{E}_{\theta_i \sim D_i}[\theta_i(a)] - p = e_a - p$, as if his θ_i^* were $(e_a, e_b, e_{\{a,b\}})$. Since this is true for any (even probabilistically chosen) outcome, i may very well act as if his true valuation were *exactly* $\theta_i^* = (e_a, e_b, e_{\{a,b\}})$.

Clarification. We could have started (and indeed originally started our paper) directly with this second model, but

1. *The Knightian model is a classical one, studied for almost a century; and*
2. *The second model appears contrived until one realizes that is logically equivalent to the Knightian one.*

The second model may appear counter-intuitive because a candidate set K_i may not be convex. In Example 2, letting $d_0 = \mathbb{E}[D_0]$ and $d_1 = \mathbb{E}[D_1]$, we have $K_i = \{d_0, d_1\}$, which is not convex if $d_0 \neq d_1$. Non convexity may be puzzling, because in a typical single-good auction we would expect that a player having a and b as two possible valuations must also have $\frac{a+b}{2}$ as another possible valuation. Let us stress that 'the possibility of holes' in K_i is *not* a restriction, nor a *speciousness* of our model. To the contrary, it is the natural sub-product of the *generality* of the Knightian setting.

To make our results stronger, when proving that a mechanism performs well, we consider all possible candidate sets, including non convex ones. When proving that a mechanism performs poorly, we only consider candidate sets that are convex.

δ-Approximate Knightian Auctions. In an auction where some of the players are Knightian, a mechanism's performance of course depends on the inaccuracy of the players' knowledge about themselves. In an auction of a single good, we say that the candidate set of a player i is δ-approximate if $\sup K_i - \inf K_i \leq \delta$, and that the auction is δ-approximate if all players have δ-approximate candidate sets. These notions naturally extend to the combinatorial auction we consider (see Section 1.1).

1.0 Prior Results and New Goals

As discussed in the Related Work Section 2, the Knightian setting has been extensively studied in decision theory. However, the only results in auctions are those in our prior work [CMZ12].

OUR PRIOR RESULTS. Notice that, although Knightian players may not be aware of their own true valuations, these valuations still exist, and the maximal social welfare (MSW) continues to be defined over them. In [CMZ12], we focused on single-good auctions and proved

- In **dominant strategies** or at **ex-post Nash equilibrium**, no mechanism —*even one allowing the players to report sets of valuations rather than a single valuation*— can guarantee social welfare greater than MSW$/n$.[2]
- In **undominated strategies**, the second-price mechanism guarantees social welfare $\geq$ MSW $-$ 2δ in single-good auctions.[3]

OUR NEW GOALS. In this paper we study the social-welfare performance of the VCG mechanism, in a Knightian setting, for the challenging *unrestricted combinatorial auctions*.[4]

[2]That is, the same guarantee offered by the trivial mechanism that, disregarding all of the players' strategies, assigns the good to a random player.

[3]In our unpublished manuscripts, we have shown that this result generalizes to all single-parameter domains using similar techniques [CMZ14d], and the Vickrey mechanism guarantees social welfare $\geq$ MSW $-$ $2m\delta$ in multi-unit auctions [CMZ14e].

[4]We acknowledge that the VCG mechanism admits computational-complexity issues [BDF$^+$10, DV11]; in this

When analyzed with Knightian players, the VCG continues to be the mechanism we all know and love, where each player i must report a single valuation, no matter how uncertain about θ_i^* he may be. In the Knightian setting, therefore, the VCG no longer is dominant-strategy. But, as recalled above, it continues to perform very well in single-good auctions.

1.1 Theorem 1: VCG Auction in Undominated Strategies

In an *(unrestricted) combinatorial auction* of n players and m goods, the set of possible allocations $\mathcal{A}$ consists of all possible partitions of $[m]$ (the set of m goods) into $1+n$ subsets $(A_0, A_1, \ldots, A_n)$, where A_0 is the (possibly empty) set of unassigned goods and A_i is the (possibly empty) set of goods assigned to player i. Given an allocation $A = (A_0, A_1, \ldots, A_n)$, player i has valuation $\theta_i^*(A_i) \in \mathbb{R}_{\geq 0}$ if $A_i \neq \varnothing$ and 0 if $A_i = \varnothing$.[5]

In a *Knightian* (unrestricted) combinatorial auction, the only information i has about the true valuation profile θ_i^* lies in K_i. Letting $K_i(S) := \{\theta_i(S)\}_{\theta_i \in K_i}$, we say that K_i is *δ-approximate* if $\sup K_i(S) - \inf K_i(S) \leq \delta$ for all non-empty $S \subseteq [m]$. We prove that,

Theorem 1 (Informal)**.** *In a δ-approximate combinatorial Knightian auction with $n \geq 2$ players and m goods, the VCG cannot, in undominated strategies, guarantee social welfare greater than* $\mathrm{MSW} - (2^{m+1} - 5)\delta$.

(The formal statement and proof of Theorem 1 can be found in Appendix 4.)

In fact, in this case we have been able to characterize UD_i, the set undominated strategies of a player i. This time, UD_i is *much larger* than K_i. Player i may choose an (almost arbitrary) constant fraction of the coordinates $\mathcal{S} \subseteq 2^{[m]}$, and deviate from $K_i(S)$ by an additive factor as large as $\Theta(2^m\delta)$ for all $S \in \mathcal{S}$. This strategy remains undominated for player i!

Perhaps more surprisingly, characterizing the undominated strategies of the VCG in unrestricted combinatorial auctions is much harder. Indeed, even *describing* the resulting set UD_i is challenging. (Indeed, we resort to geometry in order to describe it in a succinct way.)

Theorem 1 is somewhat disconcerting, if we feel that the VCG should always be the mechanism of choice for getting good social welfare, even when the players are Knightian, and even when the players are belief-free. But there are other solution concepts to consider.

1.2 Theorem 2: VCG Auctions in Regret-Minimizing Strategies

So far we have analyzed the VCG under all solution concepts traditionally used in private-value and belief-free auctions of incomplete information, assuming that the players are *utility maximizers*. We now analyze the VCG's performance in Knightian auctions in regret-minimizing strategies. The notion of a regret-minimizing strategy naturally extends to the Knightian setting. Informally, the regret of a strategy s_i of a player i is the maximum difference, taken over all possible strategy choices of i's opponents and all possible choices of θ_i in K_i, between the utility i gets by playing s_i and the utility he gets by best responding to those choices. A regret-minimizing player i chooses strategies that minimize his regret.

With respect to *pure* regret-minimizing strategies, we prove the following

paper we choose to focus on how the Knightian players rationally behave in VCG ignoring such complexity issues. It turns out this is already a very non-trivial question to tackle, not to say that in practice it is also interesting to study the VCG mechanism on selling 10 goods to 10 players, which is computationally tractable on a modern PC.

[5] All of our results for combinatorial auctions actually also hold even under a mild restriction on the players' valuation, namely, when they are *set-monotone* (or with *free disposal*): that is, $\theta_i(S) \leq \theta_i(T)$ whenever $S \subseteq T$.

Theorem 2 (Informal)**.** *In a δ-approximate combinatorial Knightian auction with n players and m goods, the VCG guarantees social welfare $\geq \text{MSW} - 2\min\{n, m\}\delta$ in pure regret-minimizing strategies.*

(We prove Theorem 2 in Section 5.)

That is, in combinatorial Knightian auctions, the performance of the VCG in (pure) regret minimizing strategies is absolutely stellar. Theorem 2 is less intuitive than it seems, because in a combinatorial, Knightian, VCG auction it is not obvious which strategies are regret-minimizing. Consider a player i who (1) happens to know that his true valuation for some subset of the good S lies in some interval $[x_S, x_S + \delta]$, and (2) chooses to play a pure, regret-minimizing strategy v_i. At first glance, it would appear that $v_i(S)$ should coincide with the center of the interval, that is, $v_i(S) = x_S + \delta/2$. In reality, however, $v_i(S)$ need not even belong to the interval $[x_S, x_S + \delta]$. Nevertheless, we prove that it cannot lie too far from the interval.

Mixed Strategies. For simplicity, Theorem 2 has been stated for pure strategies. Indeed, as shown in Appendix E.1, significant difficulties arise when dealing with mixed strategies. For instance, we must deal with the fact that a regret-minimizing mixed strategy can, in expectation and for *each* subset S, be *arbitrarily* far away from $K(S)$! However, Theorem 2 essentially continues to hold when allowing *mixed* strategies, but with a worse bound. Roughly, $\min\{n, m\}$ is replaced by n^2 (or even $n \log n$ if the valuations are set-monotone).[6]

1.3 The Meaningfulness of Theorem 2 and a Rationality Bridge Lemma

In principle, Theorem 2 or any other implementation in regret-minimizing strategies would be irrelevant, in the exact-valuation or in the Knightian setting, if at least one player is not a regret minimizer but a utility maximizer. However, we show that a separate lemma relating these two basic models of rationality *in all games* (with or without Knightian players), indicates that Theorem 2 may retain some meaningfulness. Let us explain.

- A utility-maximizing player $\mathcal{U}$ eliminates all his dominated strategies to compute his set of undominated ones, UD. Notice that $\mathcal{U}$ cannot further refine UD based on utility maximization alone. If UD consists of a single strategy s (necessarily a dominant one), then $\mathcal{U}$ of course chooses s. But:
  if UD *contains multiple strategies, which ones might* $\mathcal{U}$ *prefer?*
- A regret-minimizing player $\mathcal{R}$ eliminates all his non regret-minimizing strategies so as to compute his set of regret-minimizing strategies, RM. He might even continue this process k times, until he is satisfied or no further elimination is possible. Let us denote the final set of strategies he obtains this way by RM^k. If RM^k consists of a single strategy s, he of course chooses s. But:
  if RM^k *contains multiple strategies, which ones might* $\mathcal{R}$ *prefer?*

A possible answer is that, when he is no longer able to apply his 'favorite way of reasoning', even a die-hard utility maximizer $\mathcal{U}$ will resort to regret minimization to refine UD, and even a die-hard regret minimizer $\mathcal{R}$ will resort to utility maximization to refine RM^k. In principle, the two final sets of strategies obtained by such different refinement procedures could be vastly different. Our mentioned lemma, however, guarantees that they coincide.

[6] That is, $v_i(S) \leq v_i(T)$ for all $S \subseteq T \subseteq [m]$, all i, and all $v_i \in \Theta_i$. The interested reader can consult Appendix E for the mixed-strategy version of Theorem 2.

Abusing notation a bit, consider UD and RM also to be 'operators' acting on sets of strategies. In this case $\mathsf{UD}(\mathsf{UD}) = \mathsf{UD}$, while $\mathsf{RM}^2 \stackrel{\text{def}}{=} \mathsf{RM}(\mathsf{RM})$ may be a strict subset of RM. Then, our structural lemma can be expressed as follows.

Lemma 1 (Rationality Bridge Lemma, proved in Appendix A).

The set of strategies obtained after applying, in arbitrary order, k times the operator RM and at least once the operator UD coincides with $\mathsf{RM}^k \cap \mathsf{UD}$.

For instance, $\mathsf{RM}(\mathsf{RM}(\mathsf{UD}(\mathsf{RM}(\mathsf{RM}(\mathsf{UD}))))) = \mathsf{RM}^4(\mathsf{UD}) = \mathsf{RM}^4 \cap \mathsf{UD}$.
A formal statement and proof of the above lemma can be found in Appendix A. Here we wish just to mention the following implication for mechanism design:

For all mechanisms M and social choice correspondences f,
if M implements f in RM strategies or in UD strategies,
then M is automatically guaranteed to implement f also in $\mathsf{RM}(\mathsf{UD})$ strategies.[7]

Relative to the VCG, this guarantee implies that Theorem 2 continues to hold in $\mathsf{RM}(\mathsf{UD})$ strategies. That is, assuming that the players consider solely pure strategies,

Corollary 1. *In a δ-approximate combinatorial Knightian auction with n players and m goods, the VCG guarantees social welfare $\geq \text{MSW} - 2\min\{n, m\}\delta$ (not only when the players are regret minimizers, but also) when the players are utility maximizers who use regret only to break ties.*

(A similar corollary holds for the mentioned mixed-strategy version of Theorem 2.)

1.4 In Sum

The fact that the VCG is no longer dominant-strategy in Knightian auctions is 'no big loss'. Indeed, no dominant strategy mechanism can do better than assigning the goods at random, even in single-good auctions.

The fact that the VCG has excellent, and indeed essentially optimal, social-welfare performance in undominated strategies in multi-unit (and thus also in single-good) Knightian auctions demonstrates the wide relevance of the VCG.

The fact that the social-welfare performance of the VCG in combinatorial Knightian auctions is extremely poor in undominated strategies is just another hard fact of life. However, per the Rationality Bridging Lemma, once we assume that even die-hard utility maximizers resort to regret minimization when they are forced to break ties, then the VCG continues to be *the* mechanism of choice for good social welfare, even in the Knightian setting and in unrestricted combinatorial auctions.

In sum, as most things classical, the VCG outlives the confines in which it was conceived, and continues to be relevant in new and unforeseen settings.

1.5 Roadmap

We discuss the related work in Section 2, and provide basic definitions in Section 3.

The proof of Theorem 1 is very technically involved, so we divide it into four sections. In Section 4 we sketch a two-paged proof of a weaker form of Theorem 1 to gain intuition. In Appendix B,

[7]Indeed, for $i = 1$ the bridging lemma implies that $\mathsf{RM}(\mathsf{UD}) = \mathsf{RM} \cap \mathsf{UD} \subseteq \mathsf{RM}$. Of course, to enforce the same guarantee one could just demand that M implements f in $\mathsf{RM} \cup \mathsf{UD}$ strategies, but this is a very strong demand. Indeed $\mathsf{RM} \cup \mathsf{UD}$ could be a much larger set than $\mathsf{RM} \cap \mathsf{UD}$.

we state the stronger version of Theorem 1 that also includes the geometric characterization of the player's undominated strategies. The full proof is contained in Appendix C and D.

We provide the full proof of the pure strategy version of Theorem 2 in Section 5, and in Appendix E, we state and prove the mixed-strategy version of Theorem 2.

The proof of our structural lemma can be found in Appendix A.

2 Related Work

Models of Type Uncertainty. The Knightian model was originally proposed by Knight [Kni21] and formalized by Bewley [Bew02].

Knightian players have received much attention in *decision theory*. In particular, Aumann [Aum62], Dubra, Maccheroni and Ok [DMO04], Ok [Ok02], and Nascimento [Nas11] investigate decision with incomplete orders of preferences. Various criteria for selecting a single distribution out of a set of distributions have been studied by Danan [Dan10], Schmeidler [Sch89], Gilboa and Schmeidler [GS89]. (In fact, Bose, Ozdenoren and Pape [BOP06] and Bodoh-Creed [Bod12] use the model from [GS89] to study auctions.)

General equilibrium models with incompletely ordered preferences have been considered by Mas-Colell [Mas74], Gale and Mas-Colell [GM75], 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]

Single-player mechanisms, in the Knightian model, for the rent-extraction problem have been studied by Lopomo, Rigotti, and Shannon [LRS09], under two notions of implementation. Namely, (1) when reporting the truth is at least as good as any other strategy, and (2) when reporting the truth is not strictly eliminated in favor of another strategy.[9]

Although they are quite different from the Knightian model, a few other models of player uncertainty should be mentioned. For instance, Milgrom [Mil89], in single-good auctions, studies the revenue difference between second-price and English auctions, when the players do not exactly know their own valuations, but only that they are drawn from a *common* distribution. Sandholm [San00] presents an example of an auction (with a non quasi-linear 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. Mechanisms for scheduling, when each player knows a single distribution where his type is drawn, have been studied by Porter, Ronen, Shoham and Tennenholtz [PRST08], and by Feige and Tennenholtz [FT11]. Thompson and Leyton-Brown [TLB07] provide an extensive summary of works on Bayesian self-uncertainties.

Undominated Strategies. Implementations in undominated strategies trace back to Jackson [Jac92, JPS94]. Although being a well-known solution concept, very few positive results on mechanism design have been achieved so far. Beyond the positive example in [Jac92], Babaioff et al. [BLP06] provide an efficient mechanism for single-value multi-minded auctions, and Abreu and Matsushima [AM92] achieve perfect revenue in the complete information setting. Our prior work on the Knightian mechanism design is another example [CMZ12].

Regret-Minimizing Strategies. Regret-minimizing strategies are also known as regret-minimax strategies. The suggestion of adopting regret-minimizing (a.k.a. regret-minimax) strategies traces

[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.

[9] Notice that, not envisaging other players, these are not notions of dominance in the Knightian setting. Indeed, even in the exact-valuation setting, the notion of dominance should take into account all possible choices of strategies of the other players.

back to Savage's reading [Sav51] of the work of Wald [Wal49], and has been axiomatized by Milnor [Mil54]. The notion of regret has been treated differently in different settings. A unified axiomatic characterization of minimax regret has been recently given by Stoye [Sto11].

Mechanisms have also been studied under minimax regret. Linhart and Radner [LR89] study minimax-regret strategies in a sealed-bid mechanism for bilateral bargaining under complete information. Engelbrecht-Wiggans [Eng89] and Selten [Sel89] analyze first- and second-price sealed-bid auctions by incorporating regret for the bidders. In more general settings, minimax-regret strategies are mostly studied when a player has (Bayesian or set-theoretic) beliefs about his opponents. In particular, Hyafil and Boutilier [HB04] and Renou and Schlag [RS10] study two different notions of minimax-regret equilibrium, both coinciding with ours when players do not form beliefs about their opponents. Halpern and Pass [HP12] propose the solution concept of iterated regret minimization using beliefs.

Regret Minimizers vs. Utility Maximizers. Many empirical studies compare utility maximizers and regret minimizers, see for instance Chorus, Arentze and Timmermans [CAT09], and Hensher, Greene and Chorus [HGC11]. Recently, Engelbrecht-Wiggans and Katok [EK07] and Filiz and Ozbay [FO07] provide experimental evidence for regret in first- and second-price auctions. To the best of our knowledge, we are the first to study players who use regret for refining their sets of undominated strategies.

3 Classical and Knightian Basic Notions

Recall that, in an auction, the set of possible outcomes is $\Omega \stackrel{\text{def}}{=} \mathcal{A} \times \mathbb{R}^n_{\geq 0}$, where $\mathcal{A}$ denotes the set of all possible allocations of the good(s). If $(A, P) \in \Omega$, we refer to A, $A = (A_0, A_1, \ldots, A_n)$, as the realized allocation, to each P_i as the price charged to player i, to each A_i as the allocation of player i, and to A_0 as the unallocated good(s). A valuation θ_i of a player i is a function, from i's possible allocations to non-negative reals, mapping the empty allocation to 0. The set of all possible valuations for a player i is denoted by Θ_i, and i's true valuation by θ_i^*. We assume quasi-linear utility functions. That is, the utility function U_i of a player i maps a valuation θ_i and an outcome $\omega = (A, P)$ to $U_i(\theta_i, \omega) \stackrel{\text{def}}{=} \theta_i(A_i) - P_i$.

As already said, in a Knightian auction the only information that a player i has about θ_i^* —and the entire profile θ^*— consists of a subset $K_i \subset \Theta_i$, the candidate (valuation) set, guaranteed to contain θ_i^*. A player i has no information or belief about θ_{-i}^* or K_{-i} of his opponents. The true valuations of the players are uncorrelated.

By saying that K is a profile —respectively, a product— of candidate sets, we mean that $K = (K_1, \ldots, K_n)$ —respectively, that $K = K_1 \times \cdots \times K_n$.

Let us now clarify the specific auctions we consider.

δ-approximate Knightian Auctions. Recall that, in an (unrestricted) combinatorial auction, there are n players and m distinct goods. The set of possible allocations $\mathcal{A}$ consists of all possible partitions A of $[m]$ into $1+n$ subsets, $A = (A_0, A_1, \ldots, A_n)$, where A_0 is the (possibly empty) set of unassigned goods and A_i is the (possibly empty) set of goods assigned to player i. For each player i, $\Theta_i = \{\theta_i : 2^{[m]} \to \mathbb{R}_{\geq 0} \mid \theta_i(\varnothing) = 0\}$.

In an (unrestricted) combinatorial *Knightian* auction, a player i's candidate set K_i is a subset of the above Θ_i. If $S \subset [m]$, then we let $K_i(S) \stackrel{\text{def}}{=} \{\theta_i(S) \mid \theta_i \in K_i\}$. We say that K_i is *δ-approximate* if $\sup K_i(S) - \inf K_i(S) \leq \delta$ for all $S \subseteq [m]$.

A Knightian auction is *δ-approximate* if each candidate set K_i is δ-approximate.

(Possibly Incomplete) Preferences. In a Knightian auction, a utility-maximizing player i with candidate set K_i strictly prefers an outcome ω to an outcome ω' if and only if the following two conditions hold:

(1) $U_i(\theta_i, \omega) \geq U_i(\theta_i, \omega')$ for all $\theta_i \in K_i$ and
(2) $U_i(\theta'_i, \omega) > U_i(\theta'_i, \omega')$ for some $\theta'_i \in K_i$.

Social welfare. The social welfare of an allocation A, $\mathrm{SW}(A)$, is defined to be $\sum_i \theta^*_i(A_i)$; and the maximum social welfare, MSW, is defined to be $\max_{A \in \mathcal{A}} \mathrm{SW}(A)$. (That is, social welfare and maximum social welfare continue to be defined relative to the players' true valuations θ^*_i, whether or not the players know them exactly.)

More generally, the social welfare of an allocation A relative to a valuation profile θ, $\mathrm{SW}(\theta, A)$, is $\sum_i \theta_i(A_i)$; and the maximum social welfare relative to θ, $\mathrm{MSW}(\theta)$, is $\max_{A \in \mathcal{A}} \mathrm{SW}(\theta, A)$. Thus, $\mathrm{SW}(A) = \mathrm{SW}(\theta^*, A)$ and $\mathrm{MSW} = \mathrm{MSW}(\theta^*)$.

The VCG mechanism. In our auctions, the VCG mechanism (with any tie-breaking rule) maps a profile of valuations $\theta \in \Theta_1 \times \cdots \times \Theta_n$, to an outcome (A, P), where

$A \in \arg\max_{A \in \mathcal{A}} \mathrm{SW}(\theta, A)$ and, for each player i, $P_i = \mathrm{MSW}(\theta_{-i}) - \sum_{j \neq i} \theta(A_i)$.

General mechanisms and strategies. Every auction mechanism M considered in this paper specifies, for each player i, a set S_i. We interchangeably refer to each member of S_i as a pure *strategy/action/report* of i, and similarly, a member of $\Delta(S_i)$ a mixed strategy/action/report of i.[10] After each player i, simultaneously with his opponents, reports a strategy s_i in S_i, M maps the reported strategy profile s to an outcome $M(s) \in \Omega$. If M is probabilistic, then $M(s) \in \Delta(\Omega)$, and, for each player i, $U_i(\theta_i, M(s)) \stackrel{\text{def}}{=} \mathbb{E}_{\omega \sim M(s)}[U_i(\theta_i, \omega)]$.

Note that $S_i = \Theta_i$ in the VCG case, but in general the set S_i is arbitrary.

Knightian undominated strategies. Given a mechanism M, a pure strategy s_i of a player i with a candidate set K_i is *(weakly) undominated*,[11] in symbols $s_i \in \mathsf{UD}_i(K_i)$, if i does not have another (possibly mixed) strategy σ_i such that

(1) $\forall s_{-i}\ \forall \theta_i \in K_i \quad \mathbb{E}U_i\big(\theta_i, M(\sigma_i, s_{-i})\big) \geq U_i\big(\theta_i, M(s_i, s_{-i})\big)$, and
(2) $\exists s_{-i}\ \exists \theta_i \in K_i \quad \mathbb{E}U_i\big(\theta_i, M(\sigma_i, s_{-i})\big) > U_i\big(\theta_i, M(s_i, s_{-i})\big)$.

If K is a product/profile of candidate sets, then $\mathsf{UD}(K) \stackrel{\text{def}}{=} \mathsf{UD}_1(K_1) \times \cdots \times \mathsf{UD}_n(K_n)$.[12]

Knightian regret-minimizing strategies. Given a mechanism M, the (maximum) regret of a pure strategy s_i of a player i with candidate set K_i is

$$R_i(K_i, s_i) \stackrel{\text{def}}{=} \max_{\theta_i \in K_i} \max_{s_{-i}} \left(\max_{s'_i} U_i\big(\theta_i, M(s'_i, s_{-i})\big) - U_i\big(\theta_i, M(s_i, s_{-i})\big) \right) .$$

A pure strategy s_i is *regret-minimizing* among all pure strategies of a player i with a candidate set K_i, in symbols $s_i \in \mathsf{RM}^{\mathsf{pure}}_i(K_i)$, if $R_i(K_i, s_i) \geq R_i(K_i, s'_i)$ for all other pure strategies s'_i of i. We let $\mathsf{RM}^{\mathsf{pure}}(K) \stackrel{\text{def}}{=} \mathsf{RM}^{\mathsf{pure}}_1(K_1) \times \cdots \times \mathsf{RM}^{\mathsf{pure}}_n(K_n)$.

[10] Often, in pre-Bayesian settings, the notion of a strategy and that of an action are distinct. Indeed, a strategy s_i of a player i maps the set of all possible types of i to the set of i' possible actions/reports. But since strategies are universally quantified in all relevant definitions of this paper, we have no need to separate (and for simplicity refrain from separating) the notions of strategies and actions.

[11] This is not to be confused with the *strong dominance* that requires the inequality to be strict for all pairs (s_{-i}, θ_i). For this notion in the exact-valuation case, see for instance [FT91, LS08].

[12] As pointed out by Jackson [Jac92] in the exact-valuation case, the general notion of an undominated strategy is more complex. However, for *bounded* mechanisms, the simpler notion above coincides with the general notion, even in the Knightian setting. Since this class of mechanisms includes the VCG and all finite mechanisms, we adopt this simpler notion for this paper.

When allowing mixed strategies, the (expected) regret of a (possibly mixed) strategy σ_i of a player i with candidate set K_i is

$$R_i(K_i,\sigma_i) \stackrel{\text{def}}{=} \max_{\theta_i\in K_i}\max_{s_{-i}}\Big(\max_{s_i'} U_i\big(\theta_i, M(s_i', s_{-i})\big) - \mathbb{E}_{s_i\sim\sigma_i} U_i\big(\theta_i, M(s_i,s_{-i})\big)\Big) \ .$$

We similarly define $\mathsf{RM}_i^{\mathsf{mix}}(K_i)$ as the set of strategies of a player i that minimize regret among all mixed strategies, and let $\mathsf{RM}^{\mathsf{mix}}(K) \stackrel{\text{def}}{=} \mathsf{RM}_1^{\mathsf{mix}}(K_1)\times\cdots\times\mathsf{RM}_n^{\mathsf{mix}}(K_n)$.

4 A Weaker Version of Theorem 1

It suffices to consider the case where there are $n=2$ players, because all players other than players 1 and 2 can be made to report 0 on every subset of the goods, and thus not affect the choice of outcome. We now sketch the proof for the following slightly weaker version of Theorem 1. (We shall discuss in Appendix B the stronger statement of our theorem as well as a characterization of a player's undominated strategies.)

Theorem 1'. *In a combinatorial Knightian auction with 2 players and m goods, consider the VCG with any tie-breaking rule, then there exist products of δ-approximate candidate sets $K=K_1\times K_2$ and profiles $(v_1,v_2)\in\mathsf{UD}(K)$, such that*

$$\text{(best-case } \theta) \quad \forall\theta\in K_1\times K_2 \quad \mathrm{SW}\big(\theta,\mathsf{VCG}(v_1,v_2)\big)\le \mathrm{MSW}(\theta)-(2^m-3)\delta \tag{4.1}$$

$$\text{(worst-case } \theta) \quad \exists\theta\in K_1\times K_2 \quad \mathrm{SW}\big(\theta,\mathsf{VCG}(v_1,v_2)\big)\le \mathrm{MSW}(\theta)-(2^m-1)\delta. \tag{4.2}$$

Proof Sketch. Let $\pi_1,\dots,\pi_{2^m-1}$ be any permutation of all non-empty subsets of $[m]$ such that, whenever $j<k$, $\pi_j\not\supseteq\pi_k$.[13] We set $\pi_{2^m}\stackrel{\text{def}}{=}\pi_1$, and denote by $\overline{S}$ the complement of a subset S: that is, $\overline{S}\stackrel{\text{def}}{=}[m]\setminus S$.

We begin by choosing a highly-deviating strategy for player 1, and argue that it is undominated. Specifically, choose arbitrarily a real number x larger than δ, and then choose a candidate set K_1 and a strategy (i.e., a valuation) v_1 as follows:

$$K_1 \stackrel{\text{def}}{=} \Big\{\theta_1\in\Theta_1 \,\Big|\, \forall \text{ non-empty } S\subseteq[m],\ \theta_1(S)\in[x-\delta/2,x+\delta/2]\Big\} \text{ and}$$

$$v_1(\pi_i)\stackrel{\text{def}}{=} x+(i-1)\delta \quad \forall i\in\{1,\dots,2^m-1\} \ .$$

Note that $v_1\notin K_1$. (Indeed, $v_1(\pi_i)\in K_1(\pi_1)$ only for $i=1$.)

We now prove that the strategy v_1 is undominated. More precisely,

Claim 4.1. $v_1\in\mathsf{UD}_1(K_1)$.

Proof. We proceed by contradiction. Assume towards contradiction that v_1 is weakly dominated by a strategy $v_1'\neq v_1$. (There are two cases to consider: v_1' is pure and v_i' is mixed. For simplicity we analyze only the first one.) Assume that v_1' is pure.

(There are two cases to consider: either v_i' is a constant shift of v_i or it is not. For brevity, we analyze only the second, harder, case.) Assume that v_i' is not a constant shift of v_i. Then

$$\exists j\in\{1,\dots,2^m-1\} \quad \exists\Delta>0 \quad v_1(\pi_{j+1})-v_1(\pi_j)>\Delta>\max_{T\subseteq\pi_{j+1}} v_1'(T)-\max_{T\subseteq\pi_j} v_1'(T) \ . \tag{4.3}$$

[13] In particular, we can order the subsets of $[m]$ by increasing cardinality, and lexicographically within a given cardinality: that is, when $m=3$, $\{1\},\{2\},\{3\},\{1,2\},\{1,3\},\{2,3\},\{1,2,3\}$.

Else, that is, if for all $i \in \{1, \ldots, 2^m - 1\}$

$$v_1(\pi_{i+1}) - v_1(\pi_i) \le \max_{T \subseteq \pi_{i+1}} v_1'(T) - \max_{T \subseteq \pi_i} v_1'(T),$$

then summing up all these $2^m - 1$ inequalities we get $0 \le 0$; hence, all the inequalities are in fact tight. So there must exist some constant c such that $v_1(\pi_i) = v_1'(\pi_i) + c$ for $i \in \{1, \ldots, 2^m - 1\}$, which we have assumed not to be the case.

(There are now two more cases to consider: $j \notin \{2^m - 2, 2^m - 1\}$ and $j \in \{2^m - 2, 2^m - 1\}$. For brevity we analyze only the first, hard, one.) Assume that $j \notin \{2^m - 2, 2^m - 1\}$. In this case neither $\overline{\pi_j}$ nor $\overline{\pi_{j+1}}$ is empty.

We contradict the assumption that v_1' weakly dominates v_1 by exhibiting a valuation $\theta_1 \in K_1$ and a "witness" strategy v_2 for player 2 such that

$$U_1(\theta_1, \mathsf{VCG}(v_1, v_2)) > U_1(\theta_1, \mathsf{VCG}(v_1', v_2)) \ .$$

We define v_2 as follows. Let H be a huge number (e.g., much higher than $v_1(\pi)$ and $v_1'(\pi)$ for any subset π of the goods) and let $v_2(\overline{\pi_{j+1}}) = H - \Delta$, $v_2(\overline{\pi_j}) = H$, and $v_2(T) = 0$ for all other subsets T. (Here we rely on the combinatorial nature of the auction: we have complete freedom on how to choose the valuation v_2.)

We now argue that the allocation in the outcome $\mathsf{VCG}(v_1, v_2)$ is $(\pi_{j+1}, \overline{\pi_{j+1}})$ and player 1's price is Δ. Indeed, because H was chosen to be sufficiently large, the only outcomes we should consider are $(T, \overline{\pi_{j+1}})$ and $(T', \overline{\pi_j})$ where $T \subseteq \pi_{j+1}$ and $T' \subseteq \pi_j$. By construction π_{j+1} maximizes $v_1(T)$ among all $T \subseteq \pi_{j+1}$, and π_j maximizes $v_1(T)$ among all $T \subseteq \pi_j$; in particular, the only two possible allocations are $(\pi_j, \overline{\pi_j})$ and $(\pi_{j+1}, \overline{\pi_{j+1}})$. Because $v_1(\pi_{j+1}) - v_1(\pi_j) > \Delta = v_2(\overline{\pi_j}) - v_2(\overline{\pi_{j+1}})$, the outcome that is chosen is $(\pi_{j+1}, \overline{\pi_{j+1}})$. As for the price: player 2 is allocated $\overline{\pi_{j+1}}$ but, if player 1 did not exist, player 2 would be allocated $\overline{\pi_j}$, and gain Δ in utility; thus player 1's price is indeed Δ.

Next, we argue that the allocation in the outcome $\mathsf{VCG}(v_1', v_2)$ is $(T^*, \overline{\pi_j})$, where T^* maximizes $v_1'(T)$ among all $T \subseteq \pi_j$, and player 1's price is 0. As before, because H was chosen to be sufficiently large, the only outcomes we should consider are $(T, \overline{\pi_{j+1}})$ and $(T', \overline{\pi_j})$ where $T \subseteq \pi_{j+1}$ and $T' \subseteq \pi_j$. This time by relying on the fact that

$$v_2(\overline{\pi_j}) - v_2(\overline{\pi_{j+1}}) = \Delta > \max_{T \subseteq \pi_{j+1}} v_1'(T) - \max_{T \subseteq \pi_j} v_1'(T)$$

we deduce that the outcome is in fact $(T^*, \overline{\pi_j})$. As for the price: player 2 is allocated $\overline{\pi_j}$ and, if player 1 did not exist, player 2 would still be allocated $\overline{\pi_j}$; thus player 1's price is indeed 0.

We now define $\theta_1 \in K_1$ as follows: $\theta_1(\pi_{j+1}) = x + \delta/2$, $\theta_1(\pi_j) = x - \delta/2$, and $\theta_1(\pi)$ is arbitrarily chosen for all other subsets π. For our choices of θ_1, v_1, v_1' and v_2 we have:

$$U_1(\theta_1, \mathsf{VCG}(v_1, v_2)) = (x + \delta/2) - \Delta$$
$$U_1(\theta_1, \mathsf{VCG}(v_1', v_2)) = (x - \delta/2) - 0 \ .$$

By (4.3) and the construction of v_1, it is immediately seen that $\delta = v_1(\pi_{j+1}) - v_1(\pi_j) > \Delta$. Thus the first utility is greater than the second one, contradicting the fact that v_1' weakly dominates v_1. □

Having constructed $v_1 \in \mathsf{UD}_1(K_1)$, we continue the proof of Theorem 1' by letting:

$$v_2(S) \stackrel{\text{def}}{=} \begin{cases} (2^m - i - 1.5)\delta & \text{if } S = \overline{\pi_i} \text{ for some } i \in \{1, \ldots, 2^m - 2\} \\ x + (2^m - 2.5)\delta & \text{if } S = [m] \end{cases} ,$$

$$K_2 \stackrel{\text{def}}{=} \left\{ \theta_2 \in \Theta_2 \;\middle|\; \forall i \in \{1, \ldots, 2^m - 1\},\ \theta_2(\pi_i) \in \big[v_2(\pi_i), v_2(\pi_i) + \delta\big] \right\} \ .$$

Note that, by construction, $v_2 \in K_2$, which easily implies the following

Claim 4.2. $v_2 \in \mathsf{UD}_2(K_2)$. (For brevity we do not prove this implication.)

Having specified K_1, v_1, K_2, and v_2, all we have left is analyzing the social welfare performance.

Let us first compute the allocation of the outcome $\mathsf{VCG}(v_1, v_2)$. The only allocations to consider are $(\pi_{2^m-1}, \varnothing)$, $(\varnothing, \pi_{2^m-1})$, and $(\pi_i, \overline{\pi_i})$, for some index $i \in \{1, \dots, 2^m - 2\}$. (In principle, one may also consider allocations where some goods remain unallocated. However, since v_1 and v_2 are strictly monotone —that is, $v_j(S) < v_j(T)$ for all $S \subsetneq T$ and all $j \in \{1, 2\}$— all goods must be allocated in the outcome of $\mathsf{VCG}(v_1, v_2)$.)

Now we compare the social welfare relative to (v_1, v_2) for such allocations:

$$\begin{aligned} v_1(\pi_{2^m-1}) + v_2(\varnothing) &= (x + (2^m - 2)\delta) + 0 = x + (2^m - 2)\delta \enspace, \\ v_1(\varnothing) + v_2(\pi_{2^m-1}) &= 0 + (x + (2^m - 2.5)\delta) = x + (2^m - 2.5)\delta \enspace, \text{ and} \\ v_1(\pi_i) + v_2(\overline{\pi_i}) &= (x + (i-1)\delta) + (2^m - i - 1.5)\delta = x + (2^m - 2.5)\delta \enspace. \end{aligned}$$

Thus, in the outcome $\mathsf{VCG}(v_1, v_2)$ the allocation is $(\pi_{2^m-1}, \varnothing)$. Hence, the social welfare is

$$\mathrm{SW}\big((\theta_1, \theta_2), \mathsf{VCG}(v_1, v_2)\big) = \theta_1(\pi_{2^m-1}) \enspace.$$

On the other hand, the maximum social welfare is

$$\mathrm{MSW}(\theta_1, \theta_2) \geq \theta_2(\pi_{2^m-1}) \enspace.$$

Now notice that for all $\theta \in K$, we have

$$\mathrm{MSW}(\theta) - \mathrm{SW}\big(\theta, \mathsf{VCG}(v_1, v_2)\big) \geq \theta_2(\pi_{2^m-1}) - \theta_1(\pi_{2^m-1}) \geq (x + (2^m - 2.5)\delta) - (x + \delta/2) = (2^m - 3)\delta \enspace.$$

That is, (4.1) holds. To prove (4.2), we choose θ as follows:

$$\begin{aligned} \theta_1(\pi_i) &\stackrel{\text{def}}{=} x - \delta/2 \quad \forall\, i \in \{1, \dots, 2^m - 1\} \enspace, \\ \theta_2(\pi_i) &\stackrel{\text{def}}{=} v_2(\pi_i) + \delta \quad \forall\, i \in \{1, \dots, 2^m - 1\} \enspace. \end{aligned}$$

Now notice that

$$\mathrm{MSW}(\theta) - \mathrm{SW}\big(\theta, \mathsf{VCG}(v_1, v_2)\big) \geq \theta_2(\pi_{2^m-1}) - \theta_1(\pi_{2^m-1}) = (x + (2^m - 1.5)\delta) - (x - \delta/2) = (2^m - 1)\delta \enspace.$$

That is, (4.2) also holds. This concludes our proof sketch of the weaker version of Theorem 1. ■

5 Proof of Theorem 2

Theorem 2. *In a combinatorial Knightian auction with n players and m goods, let the VCG mechanism break ties by preferring subsets with smaller cardinalities.*[14] *Then, for all δ, all products K of δ-approximate candidate sets, all profiles $\theta \in K$, and all profiles of strategies $v \in \mathsf{RM}^{\mathsf{pure}}(K)$,*

$$\mathrm{SW}(\theta, \mathsf{VCG}(v)) \geq \mathrm{MSW}(\theta) - 2\min\{m, n\}\delta \enspace.$$

Proof. We begin by noting that, because the VCG is dominant-strategy-truthful in the exact-valuation model, the (maximum) regret of a pure strategy v_i of a player i with candidate set K_i in the VCG mechanism becomes

$$\begin{aligned} R_i(K_i, v_i) &\stackrel{\text{def}}{=} \max_{\theta_i \in K_i} \max_{v_{-i}} \Big(\max_{v'_i} U_i\big(\theta_i, \mathsf{VCG}(v'_i, v_{-i})\big) - U_i\big(\theta_i, \mathsf{VCG}(v_i, v_{-i})\big) \Big) \\ &= \max_{\theta_i \in K_i} \max_{v_{-i}} \Big(U_i\big(\theta_i, \mathsf{VCG}(\theta_i, v_{-i})\big) - U_i\big(\theta_i, \mathsf{VCG}(v_i, v_{-i})\big) \Big) \enspace, \end{aligned}$$

[14] If giving subsets A or $B \subsetneq A$ to player i provides the same social welfare, then the VCG will give B to player i.

Moreover, by the very definition of the VCG, we have

$$U_i\big(\theta_i, \mathsf{VCG}(v_i, v_{-i})\big) = \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) - \mathrm{MSW}(v_{-i}) \ .^{15}$$

Therefore in the VCG case, we can further simplify the definition of regret as follows:

$$\begin{aligned} R_i(K_i, v_i) &= \max_{\theta_i \in K_i} \max_{v_{-i}} \Big(\mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(\theta_i, v_{-i})\big) - \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) \Big) \\ &= \max_{\theta_i \in K_i} \max_{v_{-i}} \Big(\mathrm{MSW}\big(\theta_i, v_{-i}\big) - \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) \Big) \ . \end{aligned} \tag{5.1}$$

Let us adopt a notation analogous to that of the proof in [CMZ14e]. Namely, for each player i, each candidate set $K_i \subset \Theta_i$, and each subset $T \subseteq [m]$, we let

$$\begin{gathered} K_i(T) \stackrel{\text{def}}{=} \{\theta_i(T)\}_{\theta_i \in K_i}, \quad K_i^{\perp}(T) \stackrel{\text{def}}{=} \inf K_i(T), \\ K_i^{\top}(T) \stackrel{\text{def}}{=} \sup K_i(T), \quad K_i^{\text{mid}}(T) \stackrel{\text{def}}{=} (K_i^{\perp}(T) + K_i^{\top}(T))/2 \ . \end{gathered}$$

To prove Theorem 2, we rely on two intermediate claims. The first one identifies, for every player i, a strategy v_i with regret no larger than δ.

Claim 5.1. *For every player i, let $v_i^*(T) \stackrel{\text{def}}{=} K_i^{\text{mid}}(T)$ for each $T \subseteq [m]$. Then $R_i(K_i, v_i^*) \le \delta$.*

Proof of Claim 5.1. According to the first equality of (5.1), it suffices to show that

$$\forall \theta_i \in K_i \ \forall v_{-i}, \quad \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(\theta_i, v_{-i})\big) - \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i^*, v_{-i})\big) \le \delta \ .$$

Let $\omega_1 = \mathsf{VCG}(\theta_i, v_{-i})$ and $\omega_2 = \mathsf{VCG}(v_i^*, v_{-i})$.

Recall that, in a combinatorial auction, a valuation $\theta_i \in \Theta_i$ of player i maps subsets of $[m]$ to $\mathbb{R}_{\ge 0}$. For convenience, we extend θ_i to map an outcome $\omega = (A, P)$ to $\mathbb{R}_{\ge 0}$ as follows: $\theta_i(\omega) \stackrel{\text{def}}{=} \theta_i(A_i)$.

Under this notation, we have $v_i^*(\omega_2) + v_{-i}(\omega_2) \ge v_i^*(\omega_1) + v_{-i}(\omega_1)$, because the VCG maximizes social welfare relative to the strategy profile (v_i^*, v_{-i}). Using this inequality, we deduce that

$$\begin{aligned} & \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(\theta_i, v_{-i})\big) - \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i^*, v_{-i})\big) \\ = \ & \big(\theta_i(\omega_1) + v_{-i}(\omega_1)\big) - \big(\theta_i(\omega_2) + v_{-i}(\omega_2)\big) \\ = \ & \big(\theta_i(\omega_1) - \theta_i(\omega_2)\big) + \big(v_{-i}(\omega_1) - v_{-i}(\omega_2)\big) \\ \le \ & \big(\theta_i(\omega_1) - \theta_i(\omega_2)\big) + \big(v_i^*(\omega_2) - v_i^*(\omega_1)\big) \ . \end{aligned}$$

Suppose player i gets subset $T_1 \subseteq [m]$ in outcome ω_1, and subset $T_2 \subseteq [m]$ in outcome ω_2. Then

$$\begin{aligned} \big(\theta_i(\omega_1) - \theta_i(\omega_2)\big) + \big(v_i^*(\omega_2) - v_i^*(\omega_1)\big) &= \ \big(\theta_i(T_1) - v_i^*(T_1)\big) + \big(v_i^*(T_2) - \theta_i(T_2)\big) \\ &\le \ K_i^{\top}(T_1) - K_i^{\text{mid}}(T_1) + K_i^{\text{mid}}(T_2) - K_i^{\perp}(T_2) \\ &\le \ \frac{\delta}{2} + \frac{\delta}{2} = \delta \ . \qquad \square \end{aligned}$$

Let us now prove another claim.

Claim 5.2. *Let v_i be any strategy of player i such that $R_i(K_i, v_i) \le \delta$. Then:*

[15] This is because, suppose that the VCG mechanism picks an outcome $\omega = \mathsf{VCG}(v_i, v_{-i})$, allocating player i subset A_i and others A_{-i}. Then, i's price is $\mathrm{MSW}(v_{-i}) - v_{-i}(A_{-i})$ in ω. This induces a total utility of $\theta_i(A_i) + v_{-i}(A_{-i}) - \mathrm{MSW}(v_{-i}) = \mathrm{SW}((\theta_i, v_{-i}), \omega) - \mathrm{MSW}(v_{-i})$.

(a) for every $T \subseteq [m]$:

$$K_i^{\text{mid}}(T) - \max_{T' \subseteq T} v_i(T') \leq \delta - \frac{K_i^{\top}(T) - K_i^{\perp}(T)}{2} \quad , \textit{ and}$$

(b) for every $T \subseteq [m]$ such that $v_i(T) > v_i(T')$ for all $T' \subsetneq T$:

$$|v_i(T) - K_i^{\text{mid}}(T)| \leq \delta - \frac{K_i^{\top}(T) - K_i^{\perp}(T)}{2} \quad .$$

Proof. Since the case of $T = \varnothing$ is trivial, we assume below that $T \neq \varnothing$. We first prove part (a).

Suppose that (a) is not true. Then, there exists T such that

$$K_i^{\text{mid}}(T) - \max_{T' \subseteq T} v_i(T') > \delta - \frac{K_i^{\top}(T) - K_i^{\perp}(T)}{2} \quad . \tag{5.2}$$

We contradict our assumption on v_i by showing that $R_i(K_i, v_i) > \delta$.

To show $R_i(K_i, v_i) > \delta$, as per (5.1), we must find some v_{-i} and some θ_i so that

$$\text{MSW}\big(\theta_i, v_{-i}\big) - \text{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) > \delta \quad . \tag{5.3}$$

Let j be an arbitrary player other than i. We choose $\theta_i \in K_i$ such that $\theta_i(T) = K_i^{\top}(T)$,[16] and v_{-i} as follows: for every $S \subseteq [m]$

$$v_j(S) \stackrel{\text{def}}{=} \begin{cases} H & \text{if } S = \overline{T} \\ H + \varepsilon + \max_{T' \subseteq T} v_i(T') & \text{if } S = [m] \\ 0 & \text{otherwise} \end{cases} \qquad \text{and} \qquad v_k(S) \stackrel{\text{def}}{=} 0 \text{ for every } k \notin \{i, j\}.$$

Above, $\varepsilon > 0$ is some sufficiently small real number, and H is some huge real number (that is, H is much bigger than $v_i(S)$ for any subset S).[17] It then is easy to verify that the outcome $\mathsf{VCG}(v_i, v_{-i})$ allocates $\varnothing$ to player i, and $[m]$ to player j. Therefore,

$$\text{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) = \theta_i(\varnothing) + v_j([m]) = H + \varepsilon + \max_{T' \subseteq T} v_i(T') \quad .$$

On the other hand, $\text{MSW}(\theta_i, v_{-i}) \geq \theta_i(T) + v_j(\overline{T}) = K_i^{\top}(T) + H$, and therefore

$$\begin{aligned} \text{MSW}\big(\theta_i, v_{-i}\big) - \text{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) &\geq \Big(K_i^{\top}(T) + H\Big) - \Big(H + \varepsilon + \max_{T' \subseteq T} v_i(T')\Big) \\ &= K_i^{\top}(T) - \varepsilon - \max_{T' \subseteq T} v_i(T') = \frac{K_i^{\top}(T) - K_i^{\perp}(T)}{2} + K_i^{\text{mid}}(T) - \varepsilon - \max_{T' \subseteq T} v_i(T') \quad . \end{aligned}$$

Finally, since $K_i^{\text{mid}}(T) - \max_{T' \subseteq T} v_i(T')$ is strictly greater than $\delta - \frac{K_i^{\top}(T) - K_i^{\perp}(T)}{2}$, according to (5.2), there exists some sufficiently small $\varepsilon > 0$ to make $\frac{K_i^{\top}(T) - K_i^{\perp}(T)}{2} + K_i^{\text{mid}}(T) - \varepsilon - \max_{T' \subseteq T} v_i(T') > \delta$. This proves (5.3) and concludes the proof of Claim 5.2a.

We now prove part Claim 5.2b.

---

[16] Here we have implicitly assumed that $K_i^{\top}(T) = \sup K_i(T) = \max K_i(T)$, and thus we can pick $\theta_i \in K_i$ so that $\theta_i(T) = K_i^{\top}(T)$. If this is not the case, one can construct an infinite sequence $\theta_i^{(1)}, \theta_i^{(2)}, \cdots$ so that $\theta_i(T)$ approaches to $K_i^{\top}(T)$, and the rest of the proof remains unchanged.

[17] Notice that when $T = [m]$ we have $\overline{T} = \varnothing$ and one cannot assign $v_j(\varnothing)$ to be a nonzero number. In that case we can choose $H = 0$, and the rest of the proof still goes through.

One side of Claim 5.2b is easy: that is, $v_i(T) - K_i^{\text{mid}}(T) \geq -(\delta - \frac{K_i^\top(T) - K_i^\perp(T)}{2})$. Indeed, this inequality follows from $\max_{T' \subseteq T} v_i(T') = v_i(T)$ and Claim 5.2a.

To show the other side, that is, $v_i(T) - K_i^{\text{mid}}(T) \leq \delta - \frac{K_i^\top(T) - K_i^\perp(T)}{2}$, we again proceed by contradiction. Suppose there is some T such that

$$v_i(T) - K_i^{\text{mid}}(T) > \delta - \frac{K_i^\top(T) - K_i^\perp(T)}{2} \quad . \tag{5.4}$$

We contradict our assumption on v_i by showing that $R_i(K_i, v_i) > \delta$. Similarly to case (a), we need to find some v_{-i} and some θ_i so that inequality (5.3) holds.

Let j be an arbitrary player other than i. This time, we choose $\theta_i \in K_i$ such that $\theta_i(T) = K_i^\perp(T)$,[16] and choose v_{-i} as follows: for every $S \subseteq [m]$

$$v_j(S) = \begin{cases} H & \text{if } S = \overline{T} \\ H - \varepsilon + v_i(T) & \text{if } S = [m] \\ 0 & \text{otherwise} \end{cases} \qquad \text{and} \qquad v_k(S) \stackrel{\text{def}}{=} 0 \text{ for every } k \notin \{i, j\}.$$

Again, $\varepsilon > 0$ is sufficiently small, and H is huge.[17] It then is easy to verify that the outcome $\mathsf{VCG}(v_i, v_{-i})$ allocates T to player i and $\overline{T}$ to player j. Therefore,

$$\text{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) = \theta_i(T) + v_j(\overline{T}) = K_i^\perp(T) + H \quad .$$

On the other hand, $\text{MSW}(\theta_i, v_{-i}) \geq \theta_i(\varnothing) + v_j([m]) = H - \varepsilon + v_i(T)$. Therefore,

$$\begin{aligned} \text{MSW}\big(\theta_i, v_{-i}\big) - \text{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) &\geq (H - \varepsilon + v_i(T)) - (K_i^\perp(T) + H) \\ &= v_i(T) - K_i^{\text{mid}}(T) + \frac{K_i^\top(T) - K_i^\perp(T)}{2} - \varepsilon \quad . \end{aligned}$$

Finally, since $v_i(T) - K_i^{\text{mid}}(T)$ is strictly greater than $\delta - \frac{K_i^\top(T) - K_i^\perp(T)}{2}$ according to (5.4), there exists some sufficiently small $\varepsilon > 0$ to make $v_i(T) - K_i^{\text{mid}}(T) + \frac{K_i^\top(T) - K_i^\perp(T)}{2} - \varepsilon > \delta$. This proves (5.3) and concludes the proof of Claim 5.2b.

In sum, Claim 5.2 holds. □

Now we return to the proof of Theorem 2. Let $v = (v_1, \ldots, v_n) \in \mathsf{RM}^{\mathsf{pure}}(K)$ be a regret-minimizing pure strategy profile, and let $\theta \in K$ be a valuation profile.

For every player i, the strategy v_i^* (i.e., the one reporting the 'middle points') has a regret at most δ, owing to Claim 5.1. Since v_i minimizes regret among all his strategies, we immediately have $R_i(K_i, v_i) \leq R_i(v_i^*, K_i) \leq \delta$. This shows that v_i satisfies the initial hypothesis of Claim 5.2.

Now, letting $(A_0, A_1, \ldots, A_n)$ be the allocation in the outcome $\mathsf{VCG}(v_1, \ldots, v_n)$, we immediately have $v_i(A_i) \geq v_i(T')$ for any $T' \subsetneq A_i$ by the definition of the VCG. Furthermore, by our choice of the tie-breaking rule, this inequality must be strict: that is, $v_i(A_i) > v_i(T')$ for any $T' \subsetneq A_i$. Therefore, letting $T = A_i$, T satisfies the hypothesis in Claim 5.2b. Thus, we conclude that

$$\begin{aligned} \forall i \in [n], \quad |v_i(A_i) - K_i^{\text{mid}}(A_i)| \leq \delta - \frac{K_i^\top(A_i) - K_i^\perp(A_i)}{2} &\leq \delta - |\theta_i(A_i) - K_i^{\text{mid}}(A_i)| \\ &\implies |v_i(A_i) - \theta_i(A_i)| \leq \delta \quad . \end{aligned} \tag{5.5}$$

Notice that, if $A_i = \varnothing$, then $v_i(\varnothing) = \theta_i(\varnothing) = 0$.

Next, letting $(B_0, B_1, \dots, B_n)$ be the allocation that maximizes the social welfare under θ, we have

$$\sum_{i=1}^{n} v_i(A_i) \geq \sum_{i=1}^{n} \max_{T' \subseteq B_i} v_i(T') \tag{5.6}$$

because the VCG maximizes social welfare relative to $v = (v_1, \dots, v_n)$. Moreover, according to Claim 5.2a we have

$$\forall i \in [n], \quad K_i^{\mathrm{mid}}(B_i) - \max_{T' \subseteq B_i} v_i(T') \leq \delta - \frac{K_i^{\top}(B_i) - K_i^{\perp}(B_i)}{2} \leq \delta - |\theta_i(B_i) - K_i^{\mathrm{mid}}(B_i)|$$
$$\implies \theta_i(B_i) - \max_{T' \subseteq B_i} v_i(T') \leq \delta \ . \tag{5.7}$$

Also notice that, if $B_i = \varnothing$, then $\theta_i(B_i) = \max_{T' \subseteq B_i} v_i(T') = 0$.

We are now ready to compute the social welfare guarantee.

$$\begin{aligned}
\mathrm{SW}(\theta, \mathsf{VCG}(v)) = \textstyle\sum_{i=1}^{n} \theta_i(A_i) \geq \ & \textstyle\sum_{i=1}^{n} v_i(A_i) - \sum_{i \in [n], A_i \neq \varnothing} \delta && \text{(using (5.5))} \\
\geq \ & \sum_{i=1}^{n} \max_{T' \subseteq B_i} v_i(T') - \sum_{i \in [n], A_i \neq \varnothing} \delta && \text{(using (5.6))} \\
\geq \ & \textstyle\sum_{i=1}^{n} \theta_i(B_i) - \sum_{i \in [n], A_i \neq \varnothing} \delta - \sum_{i \in [n], B_i \neq \varnothing} \delta && \text{(using (5.7))} \\
\geq \ & \mathrm{MSW}(\theta) - 2\min\{n, m\}\delta \ .
\end{aligned}$$

This concludes the proof of Theorem 2. ■

Appendix

A Our Rationality Bridge Lemma

We prove the Rationality Bridge Lemma for the case envisaged in decision theory: that is, for a single player against Nature. Indeed, the generalization to n-player (strategic or pre-Bayesian) games follows as a simple corollary. This is so because, in the definitions of dominance and regret, the players' strategies are universally quantified (and so is each θ_i in K_i, if player i is Knightian), and thus can be treated as Nature's strategies.

More precisely, let $\mathcal{S}$ be a compact set of strategies of the player, and T a compact set of states of Nature.[18] We denote by U the (continuous) utility function of the player, where $U(s,t)$ is the utility under strategy $s \in \mathcal{S}$ when Nature's state is $t \in T$. Regret-minimizing strategies and undominated strategies are defined as follows:

- Given a *menu* $S \subseteq \mathcal{S}$ of strategies, the player's (maximum) regret for a strategy $s \in S$, $R_S(s)$ is the maximum difference, taken over all possible states of Nature $t \in T$, between the utility the player gets by playing s and the utility he could have gotten by best responding to t: that is, $R_S(s) \stackrel{\text{def}}{=} \max_{t \in T}\big(\max_{s^* \in S} U(s^*,t) - U(s,t)\big)$.

  The set of regret-minimizing strategies with respect to S is $\mathsf{RM}(S) \stackrel{\text{def}}{=} \arg\min_{s \in S} R_S(s)$.

- Given two strategies $s, s' \in \mathcal{S}$, s' weakly dominates s, $s' \succ s$, if

$$\forall t \in T,\ U(s',t) \geq U(s,t) \quad \text{and} \quad \exists t \in T,\ U(s',t) > U(s,t) \ .$$

  Given a menu $S \subseteq \mathcal{S}$, the player's undominated strategies consist of those that are not weakly dominated by any weakly undominated strategy.[19] That is,

$$\begin{aligned}\mathsf{UD}(S) &\stackrel{\text{def}}{=} S \setminus \{s \in S \ : \ \exists s' \in S \text{ s.t. } (s' \succ s) \wedge (\nexists s'' \in S, s'' \succ s')\} \\ &= \quad\ \{s \in S \ : \ \nexists s' \in S \text{ s.t. } (s' \succ s) \wedge (\nexists s'' \in S, s'' \succ s')\}\end{aligned}$$

We now state two simple facts that follow easily from the above definitions:

Fact A.1. *For every menu $S \subseteq \mathcal{S}$,*
(a) if $s \prec s'$ for some s and s' in S, then $R_S(s) \geq R_S(s')$, and
(b) the regrets of a strategy $s \in \mathsf{UD}(S)$ with respect to S and $\mathsf{UD}(S)$ are the same, namely:[20]

$$R_S(s) = \max_{t \in T}\Big(\max_{s^* \in S} U(s^*,t) - U(s,t)\Big) = \max_{t \in T}\Big(\max_{s^* \in \mathsf{UD}(S)} U(s^*,t) - U(s,t)\Big) = R_{\mathsf{UD}(S)}(s) \ .$$

[18] For instance, in the Knightian setting of the VCG, when analyzing a player i, $\mathcal{S}$ consists of all possible bidding strategies of player i, and T is the cartesian product of (1) all possible bidding strategy sub-profiles of i's opponents and (2) all possible true valuations of player i in his set K_i.

Both $\mathcal{S}$ and T may be infinite, and $\mathcal{S}$ may be convex in order to allow arbitrary mixed strategies to be considered.

[19] In many cases of interest (e.g., when the set of pure strategies is finite, or when the mechanism is the VCG), weakly undominated strategies coincide with undominated ones, and this is why we directly adopted that simpler notion in the main body of this paper for Knightian auctions. As argued by Jackson [Jac92], however, the above level of precision is required when handling the general case. In particular, it may happen that every pure strategy is weakly dominated by another one in an infinite chain. In such a case, all strategies are undominated but weakly dominated.

[20] The equality in the middle is because any strategy $s^* \in S \setminus \mathsf{UD}(S)$ must be weakly dominated by some $s^{**} \in S$, giving at least as good utilities as s^* for *any* $t \in T$. Therefore, such choices of s^{**} can be ignored in the inner max.

Now let us prove the following

Claim A.2. *For every menu $S \subseteq \mathcal{S}$, $\mathsf{UD}(\mathsf{RM}(S)) = \mathsf{RM}(\mathsf{UD}(S)) = \mathsf{RM}(S) \cap \mathsf{UD}(S)$.*

Proof. We divide the proof into six steps.

1. $\mathsf{RM}(\mathsf{UD}(S)) \subseteq \mathsf{RM}(S)$.

   For every $s \in \mathsf{RM}(\mathsf{UD}(S))$, we show that $s \in \mathsf{RM}(S)$ by proving that s has minimum regret among all strategies in S. Indeed:

   - For any other strategy $s' \in \mathsf{UD}(S)$, it holds that $R_{\mathsf{UD}(S)}(s) \leq R_{\mathsf{UD}(S)}(s')$. From Fact A.1b, we deduce that $R_S(s) \leq R_S(s')$.
   - For any other strategy $s' \in S \setminus \mathsf{UD}(S)$, it holds that $s' \prec s''$ for some $s'' \in \mathsf{UD}(S)$ and $R_S(s) \leq R_S(s'')$. From Fact A.1a, we deduce that $R_S(s) \leq R_S(s'') \leq R_S(s')$.

2. $\mathsf{RM}(\mathsf{UD}(S)) \subseteq \mathsf{UD}(\mathsf{RM}(S))$.

   Because $\mathsf{RM}(\mathsf{UD}(S)) \subseteq \mathsf{RM}(S)$, if there is some $s \in \mathsf{RM}(\mathsf{UD}(S))$ such that $s \notin \mathsf{UD}(\mathsf{RM}(S))$, then $s \prec s'$, for some other strategy $s' \in \mathsf{RM}(S)$ that is not weakly dominated by any other strategy in $\mathsf{RM}(S)$, by the definition of UD.

   Next, we show that s' cannot be weakly dominated by any strategy in S as well. Suppose that $s' \prec s''$ where $s'' \in S$. Then, as we have just argued, $s'' \notin \mathsf{RM}(S)$. Using Fact A.1a however, we have $R_S(s') \geq R_S(s'')$, which implies that $s'' \in \mathsf{RM}(S)$ since $s' \in \mathsf{RM}(S)$. This contradicts $s'' \notin \mathsf{RM}(S)$.

   In sum, we have shown that s is weakly dominated by $s' \in S$, and that s' is not weakly dominated by any strategy in S. This contradicts the fact that $s \in \mathsf{UD}(S)$.

3. $\mathsf{UD}(\mathsf{RM}(S)) \subseteq \mathsf{UD}(S)$.

   Suppose that there exists some $s \in \mathsf{UD}(\mathsf{RM}(S))$ that is not in $\mathsf{UD}(S)$. By the definition of $\mathsf{UD}(S)$, the strategy s must be weakly dominated by some $s' \in S$. In addition, s' is not weakly dominated by any other strategy in S. There are two cases to consider.

   - The first case is when $s' \in \mathsf{RM}(S)$. This case is impossible because $s \in \mathsf{UD}(\mathsf{RM}(S))$ implies that if s is weakly dominated by $s' \in \mathsf{RM}(S)$, then s' must also be weakly dominated, contradicting the fact that s' cannot be weakly dominated by any strategy in S.
   - The second case is when $s' \notin \mathsf{RM}(S)$. Since $s \prec s'$, by Fact A.1a, we have $R_S(s) \geq R_S(s')$. However, because $s \in \mathsf{UD}(\mathsf{RM}(S))$ implies that $s \in \mathsf{RM}(S)$, s' must be a regret minimizer with respect to S, contradicting the fact that $s' \notin \mathsf{RM}(S)$.

4. $\mathsf{UD}(\mathsf{RM}(S)) \subseteq \mathsf{RM}(\mathsf{UD}(S))$.

   Having just proved that $\mathsf{UD}(\mathsf{RM}(S)) \subseteq \mathsf{UD}(S)$, consider any strategy $s \in \mathsf{UD}(\mathsf{RM}(S))$, and suppose that $s \notin \mathsf{RM}(\mathsf{UD}(S))$. Then there exists some $s' \in \mathsf{UD}(S)$ satisfying $R_{\mathsf{UD}(S)}(s) > R_{\mathsf{UD}(S)}(s')$. This implies, by Fact A.1b, that $R_S(s) > R_S(s')$, contradicting the fact that $s \in \mathsf{RM}(S)$.

5. $\mathsf{RM}(\mathsf{UD}(S)) \subseteq \mathsf{RM}(S) \cap \mathsf{UD}(S)$.

   This step holds because $\mathsf{RM}(\mathsf{UD}(S)) = \mathsf{UD}(\mathsf{RM}(S)) \subseteq \mathsf{RM}(S)$ and $\mathsf{RM}(\mathsf{UD}(S)) \subseteq \mathsf{UD}(S)$.

   (The first equality follows from Steps 2 and 4, and the two inclusions are obvious.)

6. $\mathsf{RM}(S) \cap \mathsf{UD}(S) \subseteq \mathsf{RM}(\mathsf{UD}(S))$.

   Take any strategy $s \in \mathsf{RM}(S) \cap \mathsf{UD}(S)$, and suppose that $s \notin \mathsf{RM}(\mathsf{UD}(S))$. Then there exists some $s' \in \mathsf{UD}(S)$ satisfying $R_{\mathsf{UD}(S)}(s) > R_{\mathsf{UD}(S)}(s')$. This implies, by Fact A.1b, that $R_S(s) > R_S(s')$, contradicting the fact that $s \in \mathsf{RM}(S)$. □

Finally, the claim above easily implies the desired lemma:

Lemma 1 (Rationality Bridge Lemma)**.** *From any menu $S \subseteq \mathcal{S}$, a player who applies,* in arbitrary order*, i times the operator* RM *and at least once the operator* UD*, always obtains the same set of surviving strategies:*

$$\mathsf{RM}^i(S) \cap \mathsf{UD}(S) \ .$$

A.1 Pure vs. Mixed Strategies

So far we have been ambiguous, when discussing undominated strategies and regret-minimizing ones, about whether or not mixed strategies are allowed.

When only pure strategies are allowed, a utility maximizer plays a pure undominated strategy, and considers only pure strategies for the notion of dominance. On the other hand, a regret minimizer plays a pure strategy whose regret is minimal among his pure strategies. (For instance, Theorem 2, in the main body of this paper, is stated for pure strategies, and so is our Rationality Bridging Lemma.)

When mixed strategies are allowed, the definitions of UD and RM require more careful attention. For a regret minimizing player, the only change needed is to allow him to choose a mixed strategy that minimizes his expected regret among all his mixed strategies (see, for instance, [HB04, HP12]). Note that it is easy to construct examples in which a mixed strategy yields strictly smaller regret than any pure strategy.

It is important to realize, however, that if we allow regret minimizers to consider mixed strategies, we *should* also allow utility maximizers to consider mixed strategies. For instance, the Rationality Bridging Lemma would have difficulty equating a set of pure strategies and a set of mixed ones.

When mixed strategies are allowed, to determine whether a strategy s is weakly dominated by another strategy s', there are only two interesting cases to consider: namely, (1) s is pure and s' is mixed; and (2) both s and s' are mixed. Traditionally, the most attention has been devoted to the first case, but the second has been studied too (see for instance [CS05, RS10]). Clearly, UD can be defined in both cases, and yields a more refined set of strategies in the second case.[21] It is under this more refined case that the Rationality Bridging Lemma holds. Indeed, the just given proof of the Rationality Bridge Lemma also works when mixed strategies are allowed. This follows immediately if we explicitly define $\mathcal{S}$ to include all possibly mixed strategies of the player.

In a sense, we have nothing to lose and something to gain by adopting a more flexible definition. After all, the 'right' definitions are those yielding the 'right' theorems.

[21] Let $\mathsf{UD}^{\mathsf{pure}}$ be the set of (pure) undominated strategies in the first case, and UD be the set of (possibly mixed) undominated strategies in the second case. Then, UD is a more 'refined' notion of undominated strategies than $\mathsf{UD}^{\mathsf{pure}}$, because $\mathsf{UD}^{\mathsf{pure}} \subseteq \mathsf{UD} \subseteq \Delta(\mathsf{UD}^{\mathsf{pure}})$, i.e., $\mathsf{UD}^{\mathsf{pure}}$ coincides with the support of UD. For this reason, there is no difference in choosing between the two notions in most of the literature (see [CS05, footnote 2]).

B Theorem 1: How to Obtain a Stronger Result and a Characterization

Payoff equivalence. Two strategies s_i and s_i' are *payoff-equivalent* for player i if for any strategy sub-profile s_{-i} of i's opponents and any $\theta_i \in K_i$, player i's utilities are the same when reporting s_i or s_i'. That is, there is no difference for i to report s_i or s_i'. Given a set of strategies S_i for player i, we denote by $\widehat{S_i}$ the set that also includes every strategy of i that is payoff-equivalent to some strategy in S_i. We will use this notation to simplify our statements of the results.

Remark. Two payoff-equivalent strategies of a player i may ultimately yield different outcomes, but they are effectively the same from i's point of view. Thus a solution concept cannot be meaningful unless, when it includes a strategy profile s, it also includes all strategy profiles s' such that s_i and s_i' are payoff equivalent for a player i.

We formally state Theorem 1 as follows.

Theorem 1. *In any unrestricted combinatorial auction with n (δ-approximate Knightian) players and m goods:*

(a) For any player i with candidate set K_i, $\mathsf{UD}_i(K_i) = \widehat{\mathbf{V}(K_i)}$.
(The set of strategies $\mathbf{V}(K_i)$ is formally defined in Definition C.2, and geometrically described in Appendix B.1 below.)

(b) Even if there are only two players, there exist products of δ-approximate candidate sets $K = K_1 \times K_2$ and profiles $(v_1, v_2) \in \mathsf{UD}(K)$, such that

$$\begin{aligned} &\textit{(best-case } \theta) \quad \forall \theta \in K_1 \times K_2 \quad \mathrm{SW}\big(\theta, \mathsf{VCG}(v_1, v_2)\big) \le \mathrm{MSW}(\theta) - (2^{m+1} - 5)\delta \\ &\textit{(worst-case } \theta) \quad \exists \theta \in K_1 \times K_2 \quad \mathrm{SW}\big(\theta, \mathsf{VCG}(v_1, v_2)\big) \le \mathrm{MSW}(\theta) - (2^{m+1} - 3)\delta. \end{aligned}$$

(In Appendix C, we prove one direction of Theorem 1a: namely, $\mathsf{UD}_i(K_i) \supseteq \widehat{\mathbf{V}(K_i)}$. We shall prove $\mathsf{UD}_i(K_i) \subseteq \widehat{\mathbf{V}(K_i)}$ in the full version of the paper. In Appendix D we show how to derive Theorem 1b from Theorem 1a.)

From sketch to proof. Let us say a few words about how the sketched proof in Section 4 can be extended to a full and slightly stronger proof. The first simplification we have made is to suppose that v_1' is a pure strategy. If instead v_1' is a mixed strategy, say it equals $\sum_j p^{(j)} v_1^{(j)}$ for $\sum_j p^{(j)} = 1$ where $v_1^{(j)}$ each is a pure strategy, then the first step is to distinguish between the following three cases (at least one of them always holds):

$$\text{(a) } \exists\, j \in \{1, \dots, 2^m - 1\},\ v_1(S_{j+1}) - v_1(S_j) > \min_j \left\{ \max_{T \subseteq S_{j+1}} v_1^{(j)}(T) - \max_{T \subseteq S_j} v_1^{(j)}(T) \right\}$$

$$\text{(b) } v_1(S_1) > \min_j \left\{ \max_{T \subseteq S_1} v_1^{(j)}(T) \right\}$$

$$\text{(c) } v_1(S_1) < \max_j \left\{ \max_{T \subseteq S_1} v_1^{(j)}(T) \right\}$$

In the proof sketch above, we analyzed case (a) when v_1' happens to be a pure strategy. However, in a full proof, one has to analyze all three cases, without assuming that v_1' is pure. The analysis of each of these cases, is significantly more involved in this more general setting.

Furthermore, when analyzing case (a), we distinguished between the case $j \notin \{2^m - 2, 2^m - 1\}$ or $j \in \{2^m - 2, 2^m - 1\}$ and only analyzed the former. In the latter, the choices of "witnesses"

$\theta_1 \in K_1$ and v_2 in order to create the contradiction $U_1(\theta_1, \mathsf{VCG}(v_1, v_2)) > U_1(\theta_1, \mathsf{VCG}(v_1', v_2))$ are different. Similarly, both (b) and (c) each have a witness specially crafted for it.

Only when all of (a), (b), and (c) are fully analyzed, we can really conclude that $v_1 \in \mathsf{UD}_1(K_1)$.

Finally, even if we expect $v_2 \in \mathsf{UD}_2(K_2)$ to be true, because $v_2 \in K_2$ (and thus v_2 is not a deviating strategy), actually proving that this is the case essentially amounts to an analysis that is not much more simple than the one required to show that the highly-deviating strategy v_1 is in $\mathsf{UD}_1(K_1)$. In our full proof in Appendix D, we actually pick v_2 and K_2 more carefully (to be also highly-deviating), and doing so induces a slightly stronger result with the following social welfare upper bound:

$$\mathrm{SW}\big((\theta_1, \theta_2), \mathsf{VCG}(v_1, v_2)\big) \leq \mathrm{MSW}(\theta_1, \theta_2) - 2(2^m - 2)\delta \ .$$

B.1 Geometric Description of $\mathbf{V}(K_i)$

In this section, we just wish to provide an intuitive *description* of the set $\mathbf{V}(K_i)$, which will be formally defined in Definition C.2.

The case of two goods. We first describe $\mathbf{V}(K_i)$ in the simpler case where there are only two goods on sale (i.e., $m = 2$). In this case, the non-empty subsets of the goods are $\{1\}, \{2\}, \{1,2\}$; in particular, a valuation is a point (x, y, z) in three dimensions, and we can draw it. For the purpose of drawing, we fix the choice $K_i(\{1\}) = [6, 9]$, $K_i(\{2\}) = [8, 11]$ and $K_i(\{1,2\}) = [10, 13]$.

We begin with two simple observations:

(a) any strategy that "bids below $\min K_i(S)$ at *every* coordinate $S \subseteq [m]$" is dominated; and

(b) any strategy that "bids above $\max K_i(S)$ at *every* coordinate $S \subseteq [m]$" is dominated.

Property (a) means that a strategy v_i such that, for *every* S, $v_i(S)$ is less than $\min K_i(S)$ cannot be in $\mathbf{V}(K_i)$. That is, $\mathbf{V}(K_i)$ does not share any strategies with the following cuboid (see Figure 1a):

$$\textsc{Cuboid}_1 \stackrel{\text{def}}{=} \left\{ (x, y, z) \;\middle|\; \begin{array}{l} x < \min K_i(\{1\}) \\ y < \min K_i(\{2\}) \\ z < \min K_i(\{1,2\}) \end{array} \right\} \ .$$

Similarly, property (b) means that a strategy v_i such that, for *every* S, $v_i(S)$ is greater than $\max K_i(S)$ cannot be in $\mathbf{V}(K_i)$. That is, $\mathbf{V}(K_i)$ does not share any strategies with the following cuboid (see Figure 1b):

$$\textsc{Cuboid}_2 \stackrel{\text{def}}{=} \left\{ (x, y, z) \;\middle|\; \begin{array}{l} x > \max K_i(\{1\}) \\ y > \max K_i(\{2\}) \\ z > \max K_i(\{1,2\}) \end{array} \right\} \ .$$

Provided that a strategy v_i is neither in $\textsc{Cuboid}_1$ nor $\textsc{Cuboid}_2$ (i.e., there are S' and S'' for which $v_i(S') > \min K_i(S')$ and $v_i(S'') < \max K_i(S'')$), there can be "many ways" in which v_i could be in $\mathbf{V}(K_i)$. To express this, we need an additional definition. For valuation sets (S_1, S_2, S_3), define

$$\textsc{Cyl}(S_1, S_2, S_3) \stackrel{\text{def}}{=} \left\{ (x, y, z) \;\middle|\; \begin{array}{l} x - y \geq \min S_1 - \max S_2 \\ y - z \geq \min S_2 - \max S_3 \\ z - x \geq \min S_3 - \max S_1 \end{array} \right\} \ .$$

Note that $\textsc{Cyl}(S_1, S_2, S_3)$ is a triangular cylinder defined by three halfspaces and its axis lies on the $x = y = z$ line. For a candidate set K_i, define (see Figure 1c and 1d)

$$\begin{aligned}\textsc{Cyl}_1 &\stackrel{\text{def}}{=} \textsc{Cyl}(K_i(\{1\}), K_i(\{2\}), K_i(\{1,2\})) \\ \textsc{Cyl}_2 &\stackrel{\text{def}}{=} \text{“ } \textsc{Cyl}(K_i(\{2\}), K_i(\{1\}), K_i(\{1,2\})) \\ &\qquad\qquad \text{after the transformation } (x,y,z) \mapsto (y,x,z)\text{”}.\end{aligned}$$

Then, disregarding set boundaries, our definition of $\mathbf{V}(K_i)$ for $m = 2$ is as follows (see Figure 1e):

$$\mathbf{V}(K_i) = \textsc{Cyl}_1 \cup \textsc{Cyl}_2 - \textsc{Cuboid}_1 - \textsc{Cuboid}_2 \ .$$

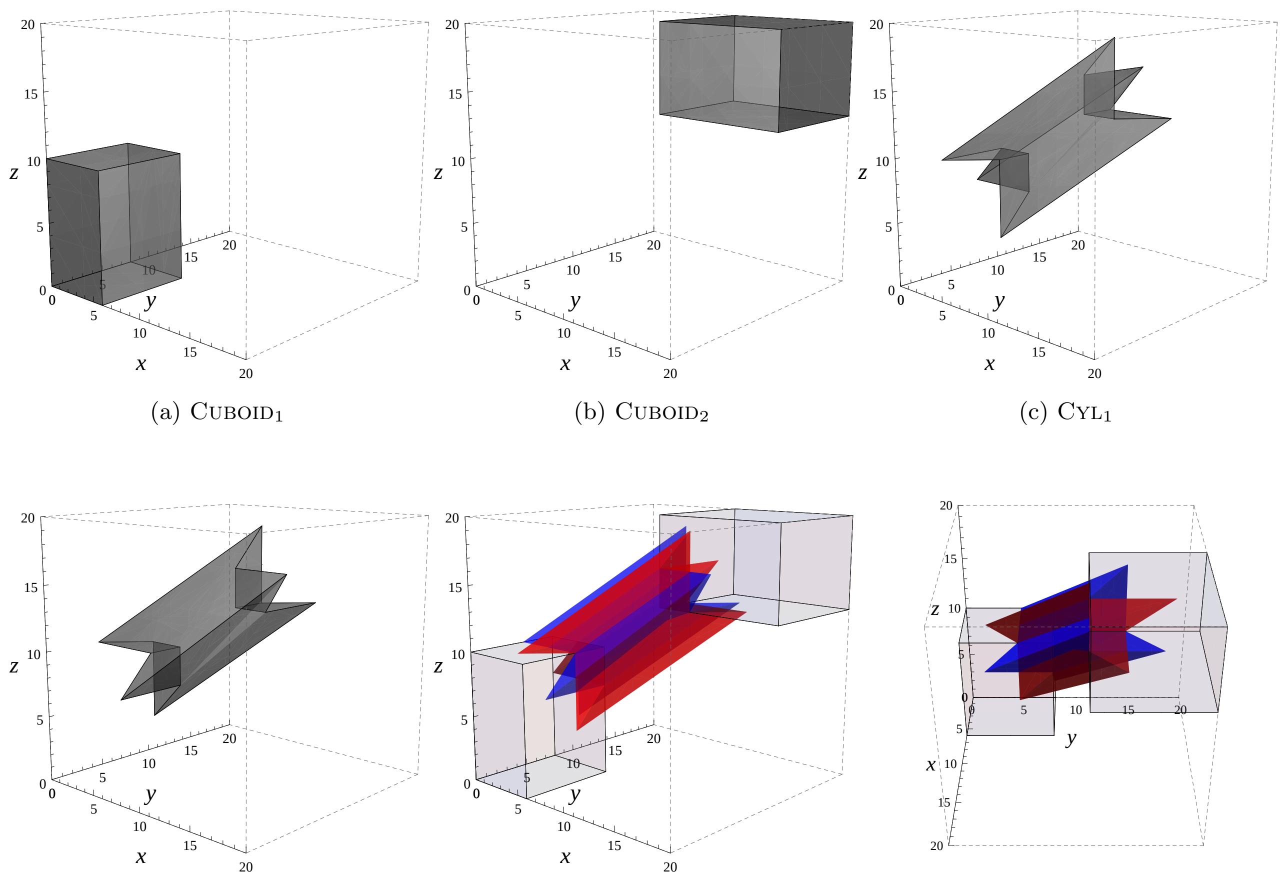

(a) $\textsc{Cuboid}_1$ (b) $\textsc{Cuboid}_2$ (c) $\textsc{Cyl}_1$

(d) $\textsc{Cyl}_2$ (e) $\mathbf{V}(K_i)$ (f)

Figure 1: Here (f) is a PDF animated rotation (if viewed under Acrobat Reader), and can also be found at `http://people.csail.mit.edu/zeyuan/knightian/vcg.gif`.

The general case. In the general case (when m need not equal 2), we can analogously define $\textsc{Cuboid}_1$ and $\textsc{Cuboid}_2$. What becomes more complicated is the “cylinder structure” of $\mathbf{V}(K_i)$. Let us explain.

When $m = 2$, there are *two* cylinders in the definition of $\mathbf{V}(K_i)$ because there are two “proper” ways of ordering all non-empty subsets of the two goods: that is $(\{1\}, \{2\}, \{1,2\})$ and $(\{2\}, \{1\}, \{1,2\})$. Thus, when $m = 2$, $\mathbf{V}(K_i)$ is the union of the two cylinders respectively obtained by indexing the three sets $K_i(\{1\})$, $K_i(\{2\})$, and $K_i(\{1,2\})$ using the two proper orderings (and minus the two cuboids).

In the general case, there are more such "proper" orderings. Concretely, we say that a relabeling π of all the non-empty subsets of $[m]$ is *proper* if $j < k$ implies that $\pi(S_j) \not\supseteq \pi(S_k)$. (Note that $\pi(S_{2^m-1}) = [m]$ is always the set of all goods.)

Analogously to the $m = 2$ case, for each vector of sets $S = (S_1, \dots, S_{2^m-1})$, we define the corresponding *fundamental cylinder* $\textsc{Cyl}(S)$. Then we consider the union of all fundamental cylinders corresponding to all vectors of sets obtained by properly relabeling $K_i = (K_i(S_1), \dots, K_i(S_{2^m-1}))$. In sum, the description of $\mathbf{V}(K_i)$ in the general case is:

$$\mathbf{V}(K_i) = \bigcup_{\substack{\text{proper} \\ \pi}} \textsc{Cyl}(\pi(K_i)) - \textsc{Cuboid}_1 - \textsc{Cuboid}_2 \ .$$

For more details see Appendix C.

C Proof of One Side of Theorem 1a

We introduce some notions before we proceed with the formal statement of the theorem. A *labeling* of all non-empty subsets of $[m]$ is a vector $\pi = (\pi_1, \dots, \pi_{2^m-1})$, where the π_i's are the $2^m - 1$ distinct non-empty subsets of $[m]$.

Definition C.1. *A labeling π of all non-empty subsets of $[m]$ is* **proper** *if $j < k \Rightarrow \pi_j \not\supseteq \pi_k$.*[22]

To make the result of our characterization clean, we assume that the candidate set K_i for the considered player i, is a cartesian product of intervals. That is, $K_i(T) = \{\theta_i(T)\}_{\theta_i \in K_i} = [a_T, b_T]$ for some $0 \le a_T \le b_T$. We denote by $K_i^\perp(T) = a_T$ the minimum point in this interval and $K_i^\top(T) = b_T$ the maximum point in this set.

Definition C.2. *For any player i with candidate set K_i, the set $\mathbf{V}(K_i)$ is the set of all strategies v_i satisfying the following two conditions:*

1. *at least one coordinate of v_i is below (resp., above) the corresponding upper (resp., lower) bound of K_i:*

$$\exists\, S' \subseteq [m], \quad v_i(S') \le K_i^\top(S') \ , \tag{C.1}$$
$$\exists\, S'' \subseteq [m], \quad v_i(S'') \ge K_i^\perp(S'') \ ; \tag{C.2}$$

2. *there exists a proper labeling π of all non-empty subsets of $[m]$ such that, letting $\pi_{2^m} \overset{\text{def}}{=} \pi_1$,*

$$\forall\, j \in \{1, \dots, 2^m - 1\}\,, \quad v_i(\pi_j) - v_i(\pi_{j+1}) \ge K_i^\perp(\pi_j) - K_i^\top(\pi_{j+1}) \ . \tag{C.3}$$

In this section we prove the harder case of Theorem 1a: its "if" side. For this side, it suffices to show that if a strategy v_i is in $\mathbf{V}(K_i)$ then it is $\mathsf{UD}_i(K_i)$. In fact, our proof assumes for simplicity that both v_i and K_i satisfy some weak monotonicity conditions. We now proceed to formally state what we are going to prove, in Lemma C.3 below.

[22] For instance, when m is equal to 3 such a permutation could be $(\{1\}, \{2\}, \{3\}, \{1,2\}, \{1,3\}, \{2,3\}, \{1,2,3\})$, or $(\{1\}, \{2\}, \{3\}, \{1,2\}, \{2,3\}, \{1,3\}, \{1,2,3\})$, and there are plenty more such permutations.

Lemma C.3 (one side of Theorem 1a)**.** *In the VCG mechanism for combinatorial auctions, no matter how ties are broken, for each player i having a weakly-monotone candidate set K_i the following holds.*[23] *If v_i is a strictly monotone strategy of i in $\mathbf{V}(K_i)$, then $v_i \in \mathsf{UD}_i(K_i)$.*

We fix a player i throughout, so we drop the subscript i everywhere. In fact, we can assume without loss of generality that $i = 1$, and that there is only another player, player 2, because all the other players can be chosen to report 0 and will thus not affect the analysis.

Assume by contradiction that a strategy v satisfying the hypothesis of the lemma is weakly dominated by some possibly mixed strategy $\{p_j, v^{(j)}\}_j$, where the probabilities p_j sum up to 1 and $v^{(j)} \neq v$ for all j. Our goal is to construct a "witness bid" $w\colon (2^{[m]} - \varnothing) \to \mathbb{R}_{\geq 0}$ for the second player and a "witness true valuation" $\theta \in K$ for the first player such that, if U is the utility function for the first player, then

$$U\big(\theta, \mathsf{VCG}(v, w)\big) > \sum_j p_j U\big(\theta, \mathsf{VCG}(v^{(j)}, w)\big) \ . \tag{C.4}$$

This will contradict the fact that the mixed strategy $\{p_j, v^{(j)}\}_j$ weakly dominates v. The construction of w and θ will be through a case analysis.

Notation.

- We call the player reporting v the first player, and the player reporting w the second player.
- We say that the allocation of $\mathsf{VCG}(v, w)$ is (S, T) if the first player receives $S \subseteq [m]$ and the second player receives $T \subseteq [m]$.
- We use $\mathrm{SW}[(S,T)] \stackrel{\text{def}}{=} v(S) + w(T)$ to denote the "apparent social welfare" of the allocation (S, T) (i.e., the social welfare when assuming that both players have the reported strategies (v, w) as their true valuations).
- Since the VCG mechanism maximizes social welfare relative to the reported strategies, we have that $\mathrm{SW}[\mathsf{VCG}(v, w)] = \max_{(S,T)}\{v(S) + w(T)\}$ where the maximization is over all $S, T \subseteq [m]$ with $S \cap T = \varnothing$.
- For notational simplicity, given a strategy v, we define its *monotonizer* $\widetilde{v}$ by $\widetilde{v}(S) \stackrel{\text{def}}{=} \max_{T \subseteq S} v(T)$.

Next, among the following inequalities, at least one cannot hold:

$$\begin{cases} v(\pi_{i+1}) - v(\pi_i) \leq \min_j \Big\{\widetilde{v^{(j)}}(\pi_{i+1}) - \widetilde{v^{(j)}}(\pi_i)\Big\}, & \forall\, i \in \{1, \dots, 2^m - 1\} \\ v(S') \leq \min_j \Big\{\widetilde{v^{(j)}}(S')\Big\} \\ v(S'') \geq \max_j \Big\{\widetilde{v^{(j)}}(S'')\Big\} \end{cases} \tag{C.5}$$

where π is any proper labeling guaranteed by the hypothesis of the lemma. Indeed, we now show that if all inequalities above hold, there must be a contradiction.

From the first inequality we deduce that, for each i and j, $v(\pi_{i+1}) - v(\pi_i) \leq \widetilde{v^{(j)}}(\pi_{i+1}) - \widetilde{v^{(j)}}(\pi_i)$; for $i \in \{1, \dots, 2^m - 1\}$, all these sum up to $0 \leq 0$. In particular, all such inequalities must be tight, so for each j, $v^{(j)}$ must be the same as v, up to a constant shift. In other words,

$$\forall\, S \subseteq [m] \text{ with } S \neq \varnothing\,, \quad v^{(j)}(S) = v(S) + c^{(j)} \text{ for some constant } c^{(j)} \ .$$

[23] A candidate set K_i is weakly monotone if $K_i^{\perp}$ and $K_i^{\top}$ are weakly monotone: for all $S, T \subseteq [m]$ with $\varnothing \subsetneq S \subseteq T$, $K_i^{\perp}(S) \leq K_i^{\perp}(T)$ and $K_i^{\top}(S) \leq K_i^{\top}(T)$. A strategy is strictly monotone if for all $S, T \subseteq [m]$ with $\varnothing \subsetneq S \subseteq T$ it holds that $v_i(S) < v_i(T)$).

Substituting the above into the second and third inequality in (C.5), we deduce that $0 \le \min_j c^{(j)}$ and $0 \ge \max_j c^{(j)}$, and therefore the $c^{(j)}$ must all be 0, contradicting the fact that $v^{(j)} \neq v$.

Therefore, one of the three kinds of inequalities in (C.5) cannot hold; we thus have three cases, depending on which kind of inequality does not hold. We now show that, for each possible case (respectively discussed in Appendix C.1, Appendix C.2, and Appendix C.3), (C.4) holds, and therefore the strategy v cannot be weakly dominated.

C.1 Case 1

Suppose that the first inequality of (C.5) does not hold for some i. For notational simplicity, assume that it does not hold for $i = 1$, i.e.,

$$v(\pi_2) - v(\pi_1) > \min_j \left\{ \widetilde{v^{(j)}}(\pi_2) - \widetilde{v^{(j)}}(\pi_1) \right\} \ .$$

We let $J = \arg\min_j \left\{ \widetilde{v^{(j)}}(\pi_2) - \widetilde{v^{(j)}}(\pi_1) \right\}$ be the set of minimizers, and let $j^* \in J$ be one of them. We can always choose some Δ such that

$$v(\pi_2) - v(\pi_1) > \Delta > \widetilde{v^{(j^*)}}(\pi_2) - \widetilde{v^{(j^*)}}(\pi_1) \ , \tag{C.6}$$

and for every $j \notin J$:

$$\widetilde{v^{(j)}}(\pi_2) - \widetilde{v^{(j)}}(\pi_1) > \Delta \ . \tag{C.7}$$

Now, we set the witness strategy of the other player to be $w(\overline{\pi_1}) = H + \Delta$, $w(\overline{\pi_2}) = H$ and $w(S) = 0$ anywhere else. Here H is some very large value. We will deal with the case when $\overline{\pi_1} = \varnothing$ or $\overline{\pi_2} = \varnothing$ later, since we cannot set the second player to have non-zero valuation on an empty set. We claim that:

Claim C.4. *If $\overline{\pi_1} \neq \varnothing$ and $\overline{\pi_2} \neq \varnothing$:*

a. The allocation of $\mathsf{VCG}(v, w)$ is $\omega = (\pi_2, \overline{\pi_2})$.

b. For all $j^* \in J$, the allocation of $\mathsf{VCG}(v^{(j^*)}, w)$ is $\omega = (T, \overline{\pi_1})$ for some $T \in \arg\max_{T \subseteq \pi_1} v^{(j^*)}(T)$ (or a probabilistic distribution over them in case of ties).

c. For all $j \notin J$, the allocation of $\mathsf{VCG}(v^{(j)}, w)$ is $\omega = (T, \overline{\pi_2})$ for some $T \in \arg\max_{T \subseteq \pi_2} v^{(j)}(T)$ (or a probabilistic distribution over them in case of ties).

Proof. For any candidate allocation (S, T) of the VCG mechanism when the second player reports w, if $T \notin \{\overline{\pi_1}, \overline{\pi_2}\}$, then $\mathrm{SW}[(S, T)]$ does not contain the big term H and is thus smaller than any $\mathrm{SW}[\omega]$ in all three cases. Therefore, we only need to consider outcomes of the form $(S, \overline{\pi_1})$ and $(S, \overline{\pi_2})$.

a. In this case, $\mathrm{SW}[\omega] = v(\pi_2) + H$. If the allocation is of the form $(S, \overline{\pi_2})$, by the strict monotonicity of v, $(\pi_2, \overline{\pi_2}) = \omega$ must be the allocation with the best social welfare. If the allocation is of the form $(S, \overline{\pi_1})$, similarly, $(\pi_1, \overline{\pi_1})$ must be the allocation with the best social welfare, however, in this case $v(\pi_1) + w(\overline{\pi_1}) = v(\pi_1) + H + \Delta < v(\pi_2) + H = \mathrm{SW}[\omega]$, using (C.6). In sum, $\omega = (\pi_2, \overline{\pi_2})$ must be the allocation of the VCG mechanism.

b. In this case, $\mathrm{SW}[\omega] = \widetilde{v^{(j^*)}}(\pi_1) + H + \Delta$. For the allocation of $(S, \overline{\pi_1})$, S must be a subset of π_1 and therefore $S \in \arg\max_{T \subseteq \pi_1} v^{(j^*)}(T)$ as desired, since the VCG mechanism is outputting an allocation with the maximum reported social welfare. For the allocation of $(S, \overline{\pi_2})$, $\mathrm{SW}[(S, \overline{\pi_2})] \le \widetilde{v^{(j^*)}}(\pi_2) + H < \widetilde{v^{(j^*)}}(\pi_1) + H + \Delta = \mathrm{SW}[\omega]$ (using (C.6)) is worse than the choice of ω. So the allocation must be of the desired form.

c. In this case, $\mathrm{SW}[\omega] = \widetilde{v^{(j)}}(\pi_2) + H$. For the allocation of $(S, \overline{\pi_1})$, we have that $\mathrm{SW}[(S, \overline{\pi_1})] \le \widetilde{v^{(j)}}(\pi_1) + H + \Delta < \widetilde{v^{(j)}}(\pi_2) + H = \mathrm{SW}[\omega]$ (using (C.7)) is worse than the choice of ω. For the allocation of $(S, \overline{\pi_2})$, S must be a subset of π_2 and therefore $S \in \arg\max_{T \subseteq \pi_2} v^{(j)}(T)$ as desired, since the VCG mechanism is outputting an allocation with the maximum reported social welfare. In sum, the allocation must be of the desired form.

□

Claim C.5. *When $\overline{\pi_1} = \varnothing$ or $\overline{\pi_2} = \varnothing$, Claim C.4 only requires the following small changes:*

a. *When $\overline{\pi_1} = \varnothing$ (i.e., $\pi_1 = [m]$), at any time $(T, \overline{\pi_1})$ is a possible allocation declared in Claim C.4, (T, R) for $R \subseteq \overline{T}$ is now also possible.*[24]

b. *When $\overline{\pi_2} = \varnothing$ (i.e., $\pi_2 = [m]$), at any time $(T, \overline{\pi_2})$ is a possible allocation declared in Claim C.4, (T, R) for $R \subseteq \overline{T}$ is now also possible.*[25]

Proof.

a. This is because, due to the (strict) monotonicity of v we have $v(\pi_1) > v(\pi_2)$ and thus (C.6) tells us that $\Delta < 0$. Instead of choosing some sufficiently large H, we can choose $H = -\Delta$. It will make sure that $w(\varnothing) = w(\overline{\pi_1}) = 0$ while $w(\overline{\pi_2}) = -\Delta > 0$. The only place that we used H being sufficiently large, is where we declare that the only possible candidate allocation for $\mathsf{VCG}(\cdot, w)$ is of the form $(S, \overline{\pi_1})$ or $(S, \overline{\pi_2})$. This is no longer true as we have to also consider (S, R) for $R \ne \overline{\pi_1}$ or $\overline{\pi_2}$. However, since $w(R) = 0$, $\mathrm{SW}[(S, R)] = \mathrm{SW}[(S, \varnothing)] = \mathrm{SW}[(S, \overline{\pi_1})]$. This means, allocation (S, R) will be possible *only if* $(S, \overline{\pi_1})$ is possible.

b. This is because, due to the weak monotonicity of $\widetilde{v^{(j^*)}}$ we have $\widetilde{v^{(j^*)}}(\pi_2) \ge \widetilde{v^{(j^*)}}(\pi_1)$ and thus (C.6) tells us that $\Delta > 0$. Instead of choosing some sufficiently large H, we can choose $H = 0$. It will make sure that $w(\varnothing) = w(\overline{\pi_2}) = 0$ while $w(\overline{\pi_1}) = \Delta > 0$. The only place that we used H being sufficiently large, is where we declare that the only possible candidate allocation for $\mathsf{VCG}(\cdot, w)$ is of the form $(S, \overline{\pi_1})$ or $(S, \overline{\pi_2})$. This is no longer true as we have to also consider (S, R) for $R \ne \overline{\pi_1}$ or $\overline{\pi_2}$. However, since $w(R) = 0$, $\mathrm{SW}[(S, R)] = \mathrm{SW}[(S, \varnothing)] = \mathrm{SW}[(S, \overline{\pi_2})]$. This means, allocation (S, R) will be possible *only if* $(S, \overline{\pi_2})$ is possible.

□

Now, we have some knowledge about what outcomes could be outputted by the VCG mechanism, on input (v, w), and on $(v^{(j)}, w)$. We now come to the final part that is to show that (C.4) holds. We first compute the utilities in all three cases:

Claim C.6. *If we choose $\theta(\pi_2) = K^{\top}(\pi_2)$ and $\theta(S) = K^{\perp}(S)$ for everything else (i.e., $S \ne \varnothing$ and $S \ne \pi_2$).*

a. *$U(\theta, \mathsf{VCG}(v, w)) = K^{\top}(\pi_2) + H - \max_S w(S)$,*

b. *$U(\theta, \mathsf{VCG}(v^{(j^*)}, w)) \le K^{\perp}(\pi_1) + H + \Delta - \max_S w(S)$ for every $j^* \in J$, and*

c. *$U(\theta, \mathsf{VCG}(v^{(j)}, w)) \le K^{\top}(\pi_2) + H - \max_S w(S)$ for every $j \notin J$.*

[24] As a consequence, Claim C.4(a) and Claim C.4(c) still hold, but Claim C.4(b) will be changed to include the possible outcomes of $\omega = (T, R)$ where T is still in $\arg\max_{T \subseteq \pi_2} v^{(j)}(T)$ but $w \subseteq \overline{T}$.

[25] As a consequence, Claim C.4(b) still holds, but Claim C.4(a) and Claim C.4(c) need small changes.

Proof.

a. We have proved in Claim C.4(a) that $(\pi_2, \overline{\pi_2})$ is the only possible allocation in this case, and therefore $U(\theta, \mathsf{VCG}(v, w)) = U(\theta, (\pi_2, \overline{\pi_2})) = K^\top(\pi_2) + w(\overline{\pi_2}) - \max_S w(S) = K^\top(\pi_2) + H - \max_S w(S)$.

b. We have proved in Claim C.4(b) that $(T, \overline{\pi_1})$ is the only possible allocation in this case, and therefore if $T \neq \pi_2$, we have $U(\theta, \mathsf{VCG}(v^{(j^*)}, w)) = K^\perp(T) + w(\overline{\pi_1}) - \max_S w(S) \leq K^\perp(\pi_1) + H + \Delta - \max_S w(S)$. (Here we used the weak monotonicity of $K^\perp$, i.e., $K^\perp(T) \leq K^\perp(\pi_1)$.)

   Otherwise, if $T = \pi_2$ (i.e., the allocation is $(\pi_2, \overline{\pi_1})$, we must have that $\pi_2 \subsetneq \pi_1$. By the (strict) monotonicity of v and (C.6), we have that $\Delta < V(\pi_2) - V(\pi_1) < 0$. In this case, since $w(\overline{\pi_1}) = H + \Delta = w(\overline{\pi_2}) + \Delta$, we know that $\mathrm{SW}[(\pi_2, \overline{\pi_2})] = \mathrm{SW}[(\pi_2, \overline{\pi_1})] - \Delta > \mathrm{SW}[(\pi_2, \overline{\pi_1})]$. This indicates that $(\pi_2, \overline{\pi_1})$ will never be a possible outcome , giving a contradiction.

c. We have proved in Claim C.4(c) that $(T, \overline{\pi_2})$ is the only possible allocation in this case, and therefore $U(\theta, \mathsf{VCG}(v^{(j^*)}, w)) \leq K^\top(T) + w(\overline{\pi_2}) - \max_S w(S) \leq K^\top(\pi_2) + w(\overline{\pi_2}) - \max_S w(S) = K^\top(\pi_2) + H - \max_S w(S)$. (Here we used the weak monotonicity of $K^\top$, i.e., $K^\top(T) \leq K^\top(\pi_2)$.)

We remark here that, in the case when $\overline{\pi_1} = \varnothing$ or $\overline{\pi_2} = \varnothing$, the allocation might also be (S, R) for some $w(R) = 0$, but one can check that the same conclusions still hold, by our choice of H.) □

Corollary C.7. *(C.4) is satisfied.*

Proof. We recall that (C.3) tells us that $v(\pi_2) - v(\pi_1) \leq K^\top(\pi_2) - K^\perp(\pi_1)$, but we have $v(\pi_2) - v(\pi_1) > \Delta$ in (C.6). This tells us that $K^\top(\pi_2) > K^\perp(\pi_1) + \Delta$.

Now, for every $j^* \in J$,

$$U(\theta, \mathsf{VCG}(v, w)) = K^\top(\pi_2) + H - \max_S w(S) > K^\perp(\pi_1) + H + \Delta - \max_S w(S) \geq U(\theta, \mathsf{VCG}(v^{(j^*)}, w))$$

while for every $j \notin J$,

$$U(\theta, \mathsf{VCG}(v, w)) = K^\top(\pi_2) + H - \max_S w(S) \geq U(\theta, \mathsf{VCG}(v^{(j)}, w))$$

The combination of them immediately implies (C.4) □

We recall that (C.4) gives a contradiction and says that v is an undominated strategy, and this ends the proof of Lemma C.3, for Case 1.

C.2 Case 2

Suppose that the second inequality of (C.5) does not hold, that is, $v(S') > \min_j\{v^{(j)}(S')\}$. Similarly as in Case 1, we let $J = \arg\min_j \left\{\widetilde{v^{(j)}}(S')\right\}$ be the set of minimizers, and let $j^* \in J$ be one of them. We can always choose some Δ such that

$$v(S') > \Delta > \widetilde{v^{(j^*)}}(S') \ , \tag{C.8}$$

and for every $j \notin J$:

$$\widetilde{v^{(j)}}(S') > \Delta \ . \tag{C.9}$$

Now, consider the following witness player, with $w(\overline{S'}) = H$ and $w([m]) = H + \Delta$, and $w(S) = 0$ everywhere else. Notice that unlike Case 1, $\Delta > 0$ is always positive. We also let H be sufficiently large when $\overline{S'} \neq \varnothing$. We choose $H = 0$ if $\overline{S'} = \varnothing$.

Claim C.8 (A variant of Claim C.4). *If $\overline{S'} \neq \varnothing$,*

a. *The allocation of* $\mathsf{VCG}(v, w)$ *is* $\omega = (S', \overline{S'})$

b. *For all* $j^* \in J$, *the allocation of* $\mathsf{VCG}(v^{(j^*)}, w)$ *is* $\omega = (\varnothing, [m])$.

c. *For all* $j \notin J$, *the allocation of* $\mathsf{VCG}(v^{(j)}, w)$ *is* $\omega = (T, \overline{S'})$, *where* $T \in \arg\max_{T \subseteq S'} v^{(j)}(T)$ *(or a probabilistic distribution over them in case of ties).*

Proof. For any candidate allocation (S, T) of the VCG mechanism when the second player reports w, if $T \notin \{\overline{S'}, [m]\}$, then $\mathrm{SW}[(S, T)]$ does not contain the big term H and is thus smaller than any $\mathrm{SW}[\omega]$ in all three cases. Therefore, we only need to consider outcomes of the form $(S, \overline{S'})$ and $(\varnothing, [m])$.

a. In this case, $\mathrm{SW}[\omega] = v(S') + H$. If the allocation is of the form $(S, \overline{S'})$, by the strict monotonicity of v, $(S', \overline{S'}) = \omega$ must be the allocation with the best social welfare. If the allocation is $(\varnothing, [m])$ its social welfare $\mathrm{SW}[(\varnothing, [m])] = \Delta + H < v(S') + H = \mathrm{SW}[\omega]$, using (C.8). In sum, $\omega = (S', \overline{S'})$ must be the allocation of the VCG mechanism.

b. In this case, $\mathrm{SW}[\omega] = H + \Delta$. For the allocation of the form $(S, \overline{S'})$, $\mathrm{SW}[(S, \overline{S'})] \leq \widetilde{v^{(j^*)}}(S) + H < H + \Delta = \mathrm{SW}[\omega]$ (using (C.8)) is worse than the choice of ω.

c. In this case, $\mathrm{SW}[\omega] = \widetilde{v^{(j)}}(S') + H$. For the allocation of $(\varnothing, [m])$, we have that $\mathrm{SW}[(\varnothing, [m])] = H + \Delta < \widetilde{v^{(j)}}(S') + H = \mathrm{SW}[\omega]$ (using (C.9)) is worse than the choice of ω. For the allocation of the form $(S, \overline{S'})$, S must be a subset of S' and therefore $S \in \arg\max_{T \subseteq S'} v^{(j)}(T)$ as desired, since the VCG mechanism is outputting an allocation with the maximum reported social welfare. In sum, the allocation must be of the desired form.

□

Claim C.9 (A variant of Claim C.5). *When $\overline{S'} = \varnothing$ (i.e., $S' = [m]$), Claim C.8 only requires the following small changes:*

> *at any time $(T, \overline{S'})$ is a possible allocation declared in Claim C.8, (T, R) for $R \subseteq \overline{T}$ is now also possible.*[26]

Proof. Recall that, instead of choosing some sufficiently large H, we choose $H = 0$ in this case. The only place that we used H being sufficiently large, is where we declare that the only possible candidate allocation for $\mathsf{VCG}(\cdot, w)$ is of the form $S, \overline{S'})$ or $(\varnothing, [m])$. This is no longer true as we have to also consider (S, R) for $R \neq \overline{S'}$ or $[m]$. However, since $w(R) = 0$, $\mathrm{SW}[(S, R)] = \mathrm{SW}[(S, \varnothing)] = \mathrm{SW}[(S, \overline{S'})]$. This means, allocation (S, R) will be possible *only if* $(S, \overline{S'})$ is possible. □

Now, we have some knowledge about what outcomes could be outputted by the VCG mechanism, on input (v, w) and on $(v^{(j)}, w)$. We now come to the final part that is to show that (C.4) holds. We first compute the utilities in all three cases:

Claim C.10 (A variant of Claim C.6). *If we choose $\theta(S) = K^\top(S)$ for everything non-empty S:*

a. $U(\theta, \mathsf{VCG}(v, w)) = K^\top(S') + H - \max_S w(S)$,

b. $U(\theta, \mathsf{VCG}(v^{(j^*)}, w)) = H + \Delta - \max_S w(S)$ *for every* $j^* \in J$, *and*

[26] As a consequence, Claim C.8(b) still holds, but Claim C.8(a) and Claim C.8(c) need small changes.

c. $U(\theta, \mathsf{VCG}(v^{(j)}, w)) \le K^\top(S') + H - \max_S w(S)$ *for every* $j \notin J$.

Proof.

a. We have proved in Claim C.8(a) that $(S', \overline{S'})$ is the only possible allocation in this case, and therefore $U(\theta, \mathsf{VCG}(v, w)) = U(\theta, (S', \overline{S'})) = K^\top(S') + w(\overline{S'}) - \max_S w(S) = K^\top(S') + H - \max_S w(S)$.

   (In the case when $\overline{S'} = \varnothing$, the allocation might also be (S', R) for some $w(R) = 0$, and since we have chosen $H = 0$ this utility equation still holds.)

b. We have proved in Claim C.8(b) that $(\varnothing, [m])$ is the only possible allocation in this case, and therefore $U(\theta, \mathsf{VCG}(v^{(j^*)}, w)) = 0 + w([m]) - \max_S w(S) = H + \Delta - \max_S w(S)$.

c. We have proved in Claim C.8(c) that $(T, \overline{S'})$ is the only possible allocation in this case, and therefore $U(\theta, \mathsf{VCG}(v^{(j)}, w)) \le K^\top(T) + w(\overline{S'}) - \max_S w(S) \le K^\top(S') + w(\overline{S'}) - \max_S w(S) = K^\top(S') + H - \max_S w(S)$.

   (Here we used the weak monotonicity of $K^\top$, i.e., $K^\top(T) \le K^\top(S')$. In the case when $\overline{S'} = \varnothing$, the allocation might also be (T, R) for some $w(R) = 0$, and since we have chosen $H = 0$ this utility equation still holds.)

□

Corollary C.11. *(C.4) is satisfied.*

Proof. We recall that (C.1) and (C.8) tell us that $\Delta < v(S') \le K^\top(S')$. Now, for every $j^* \in J$,

$$U(\theta, \mathsf{VCG}(v, w)) = K^\top(S') + H - \max_S w(S) > H + \Delta - \max_S w(S) = U(\theta, \mathsf{VCG}(v^{(j^*)}, w))$$

while for every $j \notin J$,

$$U(\theta, \mathsf{VCG}(v, w)) = K^\top(S') + H - \max_S w(S) \ge U(\theta, \mathsf{VCG}(v^{(j)}, w))$$

The combination of them immediately implies (C.4) □

We recall that (C.4) gives a contradiction and says that v is an undominated strategy, and this ends the proof of Lemma C.3, for Case 2.

C.3 Case 3

Suppose that the second inequality of (C.5) does not hold, that is, $v(S'') < \max_j \{v^{(j)}(S'')\}$. Similarly as in Cases 1 and 2, we let $J = \arg\max_j \left\{ \widetilde{v^{(j)}}(S'') \right\}$ be the set of maximizers, and let $j^* \in J$ be one of them. We can always choose some Δ such that

$$v(S'') < \Delta < \widetilde{v^{(j^*)}}(S'') \ , \tag{C.10}$$

and for every $j \notin J$:

$$\widetilde{v^{(j)}}(S'') < \Delta \ . \tag{C.11}$$

Now, consider the following witness player, with $w(\overline{S''}) = H$ and $w([m]) = H + \Delta$, and $w(S) = 0$ everywhere else. Notice that unlike Case 1, $\Delta > 0$ is always positive. We also let H be sufficiently large when $\overline{S''} \ne \varnothing$. We choose $H = 0$ if $\overline{S''} = \varnothing$.

Claim C.12 (A variant of Claim C.4). *If $\overline{S''} \neq \varnothing$,*

a. *The allocation of* $\mathsf{VCG}(v, w)$ *is* $\omega = (\varnothing, [m])$.

b. *For all* $j^* \in J$*, the allocation of* $\mathsf{VCG}(v^{(j^*)}, w)$ *is* $\omega = (T, \overline{S''})$*, where* $T \in \arg\max_{T \subseteq S''} v^{(j)}(T)$ *(or a probabilistic distribution over them in case of ties).*

c. *For all* $j \notin J$*, the allocation of* $\mathsf{VCG}(v^{(j)}, w)$ *is* $\omega = (\varnothing, [m])$.

Proof. For any candidate allocation (S, T) of the VCG mechanism when the second player reports w, if $T \notin \{\overline{S''}, [m]\}$, then $\mathrm{SW}[(S, T)]$ does not contain the big term H and is thus smaller than any $\mathrm{SW}[\omega]$ in all three cases. Therefore, we only need to consider outcomes of the form $(S, \overline{S''})$ and $(\varnothing, [m])$.

a. In this case, $\mathrm{SW}[\omega] = H + \Delta$. If the allocation is of the form $(S, \overline{S''})$, by the strict monotonicity of v, $(S'', \overline{S''}) = \omega$ must be the allocation with the best social welfare. However, its social welfare $\mathrm{SW}[(S'', \overline{S''})] = v(S'') + H < H + \Delta = \mathrm{SW}[\omega]$, using (C.10). In sum, $(\varnothing, [m])$ must be the allocation of the VCG mechanism.

b. In this case, $\mathrm{SW}[\omega] = \widetilde{v^{(j^*)}}(S'') + H$. For the allocation of $(\varnothing, [m])$, we have that $\mathrm{SW}[(\varnothing, [m])] = H + \Delta < \widetilde{v^{(j^*)}}(S'') + H = \mathrm{SW}[\omega]$ (using (C.10)) is worse than the choice of ω. For the allocation of the form $(S, \overline{S''})$, S must be a subset of S'' and therefore $S \in \arg\max_{T \subseteq S''} v^{(j^*)}(T)$ as desired, since the VCG mechanism is outputting an allocation with the maximum reported social welfare. In sum, the allocation must be of the desired form.

c. In this case, $\mathrm{SW}[\omega] = H + \Delta$. For the allocation of the form $(S, \overline{S''})$, $\mathrm{SW}[(S, \overline{S''})] \leq \widetilde{v^{(j)}}(S) + H < H + \Delta = \mathrm{SW}[\omega]$ (using (C.11)) is worse than the choice of ω.

□

Claim C.13 (A variant of Claim C.5). *When $\overline{S''} = \varnothing$ (i.e., $S'' = [m]$), Claim C.12 only requires the following small changes:*

> *at any time $(T, \overline{S''})$ is a possible allocation declared in Claim C.12, (T, R) for $R \subseteq \overline{T}$ is now also possible.*[27]

Proof. Recall that, instead of choosing some sufficiently large H, we choose $H = 0$ in this case. The only place that we used H being sufficiently large, is where we declare that the only possible candidate allocation for $\mathsf{VCG}(\cdot, w)$ is of the form $(S, \overline{S''})$ or $(\varnothing, [m])$. This is no longer true as we have to also consider (S, R) for $R \neq \overline{S''}$ or $[m]$. However, since $w(R) = 0$, $\mathrm{SW}[(S, R)] = \mathrm{SW}[(S, \varnothing)] = \mathrm{SW}[(S, \overline{S''})]$. This means, allocation (S, R) will be possible *only if* $(S, \overline{S''})$ is possible. □

Now, we have some knowledge about what outcomes could be outputted by the VCG mechanism, on input (v, w) and on $(v^{(j)}, w)$. We now come to the final part that is to show that (C.4) holds. We first compute the utilities in all three cases:

Claim C.14 (A variant of Claim C.6). *If we choose $\theta(S) = K^{\perp}(S)$ for all non-empty S:*

a. $U(\theta, \mathsf{VCG}(v, w)) = H + \Delta - \max_S w(S)$,

b. $U(\theta, \mathsf{VCG}(v^{(j^*)}, w)) \leq H + K^{\perp}(S'') - \max_S w(S)$ *for every* $j^* \in J$*, and*

[27] As a consequence, Claim C.12(a) and Claim C.12(c) still hold, but Claim C.12(b) needs small changes.

c. $U(\theta, \mathsf{VCG}(v^{(j)}, w)) = \Delta + H - \max_S w(S)$ *for every* $j \notin J$.

Proof.

a. We have proved in Claim C.12(a) that $(\varnothing, [m])$ is the only possible allocation in this case, and therefore $U(\theta, \mathsf{VCG}(v, w)) = U(\theta, (\varnothing, [m])) = 0 + w(\overline{S''}) - \max_S w(S) = H + \Delta - \max_S w(S)$.

b. We have proved in Claim C.12(b) that $(T, \overline{S''})$ is the only possible allocation in this case, and therefore $U(\theta, \mathsf{VCG}(v^{(j^*)}, w)) \leq K^{\perp}(T) + w(\overline{S''}) - \max_S w(S) \leq K^{\perp}(S'') + w(\overline{S''}) - \max_S w(S) = K^{\perp}(S'') + H - \max_S w(S)$.

   (Here we used the weak monotonicity of $K^{\perp}$, i.e., $K^{\perp}(T) \leq K^{\perp}(S'')$. In the case when $\overline{S''} = \varnothing$, the allocation might also be (T, R) for some $w(R) = 0$, and since we have chosen $H = 0$ this utility equation still holds.)

c. We have proved in Claim C.12(c) that $(\varnothing, [m])$ is the only possible allocation in this case, and therefore $U(\theta, \mathsf{VCG}(v^{(j)}, w)) = 0 + w([m]) - \max_S w(S) = H + \Delta - \max_S w(S)$.

□

Corollary C.15. *(C.4) is satisfied.*

Proof. We recall that (C.2) and (C.10) tell us that $\Delta > v(S'') \geq K^{\perp}(S'')$. Now, for every $j^* \in J$,

$$U(\theta, \mathsf{VCG}(v, w)) = H + \Delta - \max_S w(S) > H + K^{\perp}(S'') - \max_S w(S) = U(\theta, \mathsf{VCG}(v^{(j^*)}, w))$$

while for every $j \notin J$,

$$U(\theta, \mathsf{VCG}(v, w)) = H + \Delta - \max_S w(S) = U(\theta, \mathsf{VCG}(v^{(j)}, w))$$

The combination of them immediately implies (C.4) □

We recall that (C.4) gives a contradiction and says that v is an undominated strategy, and this ends the proof of Lemma C.3, for Case 3.

D Proof of Theorem 1b

Theorem 1b (restated)**.** *In a combinatorial Knightian auction with* 2 *players and* m *goods, consider the VCG with any tie-breaking rule, then there exist products of* δ*-approximate candidate sets* $K = K_1 \times K_2$ *and profiles* $(v_1, v_2) \in \mathsf{UD}(K)$*, such that*

$$\text{(best-case } \theta\text{)} \quad \forall \theta \in K_1 \times K_2 \quad \mathrm{SW}\big(\theta, \mathsf{VCG}(v_1, v_2)\big) \leq \mathrm{MSW}(\theta) - (2^{m+1} - 5)\delta \tag{D.1}$$

$$\text{(worst-case } \theta\text{)} \quad \exists \theta \in K_1 \times K_2 \quad \mathrm{SW}\big(\theta, \mathsf{VCG}(v_1, v_2)\big) \leq \mathrm{MSW}(\theta) - (2^{m+1} - 3)\delta. \tag{D.2}$$

We prove the theorem in two steps.

Step 1 (Appendix D.1). We construct a candidate hard instance for the VCG mechanism, by specifying two candidate sets K_1 and K_2 and two corresponding undominated strategies v_1 and v_2, for player 1 and player 2 respectively. To show that indeed $v_1 \in \mathsf{UD}_1(K_1)$ and $v_2 \in \mathsf{UD}_2(K_2)$, we prove that our choices of v_1 and v_2 do satisfy the requirements given in Lemma C.3.

Step 2 (Appendix D.2). We show that if player 1 has candidate set K_1 and reports v_1, and player 2 has candidate set K_2 and reports v_2 (while other players report 0), the fraction of the maximum social welfare that is guaranteed is at most the value stated in the theorem.

D.1 Construction of The Hard Instance

We construct two candidate sets K_1 and K_2 and two strategies v_1 and v_2 where, for $i = 1, 2$, K_i and v_i together satisfy the hypothesis of Lemma C.3; we deduce that for our choices it holds that $v_1 \in \mathsf{UD}_1(K_1)$ and $v_2 \in \mathsf{UD}_2(K_2)$. These choices form our candidate hard instance for the VCG mechanism. (We carry out the social welfare analysis in Section D.2.)

Fix any labeling π over all $2^m - 1$ non-empty subsets of $[m]$ such that:

1. if $i < j$, then $\pi_i \not\supseteq \pi_j$ (i.e., π is proper, cf. Definition C.1);
2. $\pi_i = \overline{\pi_{2^m-1-i}}$; and
3. $\pi_{2^m-1} = [m]$.

For instance, when $m = 3$ we can let $\pi = (\{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\})$. It is a simple exercise to prove that such a π exists for any $m \geq 2$.

Also fix any positive constant x (which should be thought of as a large constant).

We begin by choosing K_1 and v_1 (depending on π and x), and showing that $v_1 \in \mathsf{UD}_1(K_1)$:

Claim D.1. *Choose:*

- K_1 *to be such that* $K_1(\pi_i) = [x - \delta/2, x + \delta/2]$ *for all* $i \in \{1, \ldots, 2^m - 1\}$.
- v_1 *to be such that* $v_1(\pi_i) = x + (i - 1)\delta$ *for all* $i \in \{1, \ldots, 2^m - 1\}$.

(See Figure 2a.) Then $v_1 \in \mathsf{UD}_1(K_1)$.

Proof. It suffices to verify that the assumptions in Lemma C.3 hold. Indeed, $K_1^\perp$ and $K_1^\top$ are both weakly monotone because they are constant; v_1 is strictly monotonic since $v_1(\pi_i) < v_1(\pi_j)$ if $i < j$. If we choose $S' = S'' = \pi_1$, we definitely have $v_1(S') = x \leq x + \delta/2 = K_1^\top(S')$ and $v_1(S'') = x \geq x - \delta/2 = K_1^\perp(S'')$. Finally, we are left to verify (C.3), and we need a "witness labeling" for that. We simply choose π to be this labeling for which we have:

$$\forall\, i \in \{1, \ldots, 2^m - 2\}, \qquad v_1(\pi_i) - v_1(\pi_{i+1}) = -\delta = K_1^\perp(\pi_i) - K_1^\top(\pi_{i+1}) \ ,$$

and for $i = 2^m - 1$,

$$v_1(\pi_{2^m-1}) - v_1(\pi_1) > 0 > -\delta = K_1^\perp(\pi_{2^m-1}) - K_1^\top(\pi_1) \ .$$

This ends the proof that $v_1 \in \mathsf{UD}_1(K_1)$. □

Next, fixing any positive constant ε (which should be thought of as a small constant), we choose K_2 and v_2 (depending on π, x, and ε), and show that $v_2 \in \mathsf{UD}_2(K_2)$:

Claim D.2. *Choose:*

- K_2 *to be such that*
$$K_2(\pi_i) = \big[(2i - 1)\delta - \varepsilon, 2i\delta - \varepsilon\big]$$
for all $i \in \{1, \ldots, 2^m - 2\}$*, and* $K_2(\pi_{2^m-1})$ *to be*
$$K_2(\pi_{2^m-1}) = \big[x + 2(2^m - 2)\delta - \varepsilon, x + (2(2^m - 2) + 1)\delta - \varepsilon\big] \ .$$
- v_2 *to be such that* $v_2(\pi_i) = i\delta - \varepsilon$ *for all* $i \in \{1, \ldots, 2^m - 2\}$*, and* $v_2(\pi_{2^m-1}) = x + (2^m - 2)\delta - \varepsilon$.

(See Figure 2b.) Then $v_2 \in \mathsf{UD}_2(K_2)$ *owing to Lemma C.3.*

Proof. First, for sufficiently small ε, $K_2^\perp(\pi_i)$ and $K_2^\top(\pi_i)$ are both positive. Once again it suffices to verify that the assumptions in Lemma C.3 hold. Indeed, $K_2^\perp$, $K_2^\top$, and v_2 are all strictly monotonic:

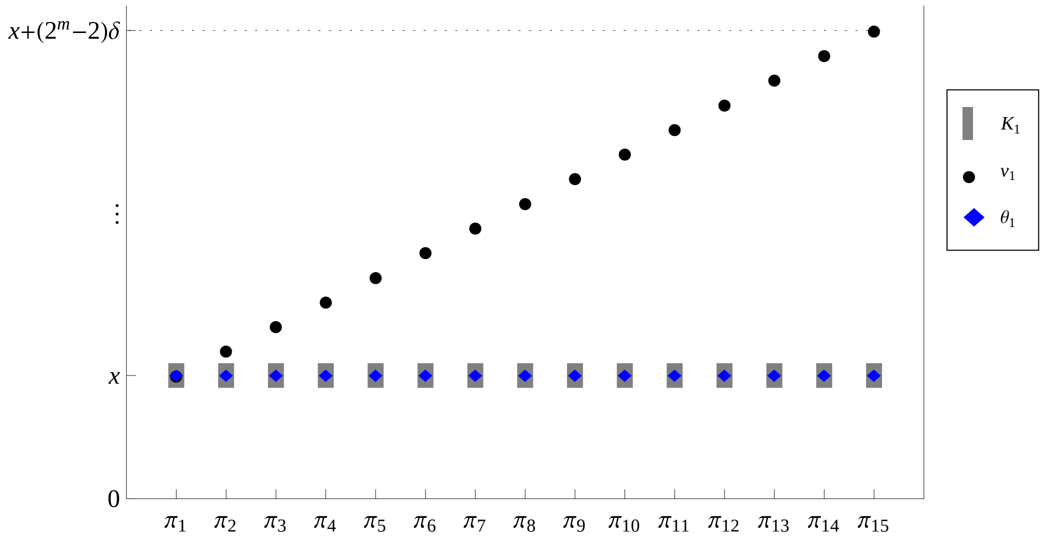

(a) K_1 and v_1 with $v_1 \in \mathsf{UD}_1(K_1)$, and $\theta_1 \in K_1$

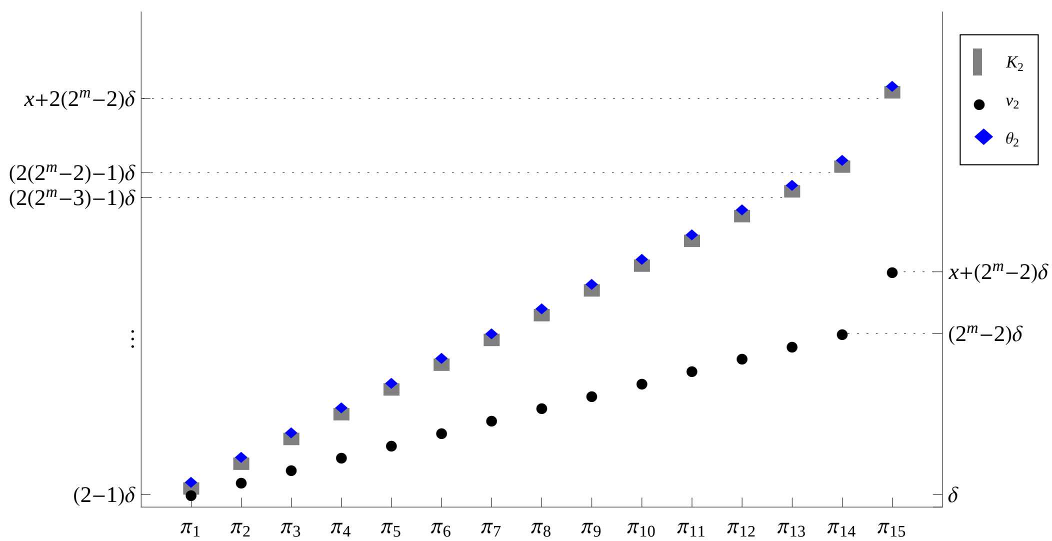

(b) K_2 and v_2 with $v_2 \in \mathsf{UD}_2(K_2)$, and $\theta_2 \in K_2$

Figure 2: The two hard instances constructed in Section D.1 and the choice of true valuation made in Section D.2, for the special case of $m = 4$.

- $K_2^{\perp}(\pi_i) < K_2^{\perp}(\pi_j)$ for $i < j$,
- $K_2^{\top}(\pi_i) < K_2^{\top}(\pi_j)$ for $i < j$, and
- $v_2(\pi_i) < v_2(\pi_j)$ if $i < j$.

If we choose $S' = S'' = \pi_1$, we have $v_2(S') = \delta - \varepsilon < 2\delta - \varepsilon = K_2^{\top}(S')$ and $v_2(S'') = \delta - \varepsilon = K_2^{\perp}(S'')$. We are left to verify (C.3), and we need a "witness labeling" for that. We now choose the labeling that is the "reverse" of π, i.e., we let $\pi'_i = \pi_{2^m - i}$; for this choice of π' we have:

- for $2 \le i \le 2^m - 2$ and let $j = 2^m - 1 - i \in \{1, 2, \dots, 2^m - 3\}$:

$$\begin{aligned} & v_2(\pi'_i) - v_2(\pi'_{i+1}) = v_2(\pi_{j+1}) - v_2(\pi_j) = \delta \\ = \; & \big((2(j+1) - 1)\delta - \varepsilon\big) - \big(2j\delta - \varepsilon\big) \\ = \; & K_2^{\perp}(\pi_{j+1}) - K_2^{\top}(\pi_j) = K_2^{\perp}(\pi'_i) - K_2^{\top}(\pi'_{i+1}) \ , \end{aligned}$$

- for $i = 1$:

$$\begin{aligned} & v_2(\pi'_i) - v_2(\pi'_{i+1}) = v_2(\pi_{2^m-1}) - v_2(\pi_{2^m-2}) = x \\ = \; & K_2^{\perp}(\pi_{2^m-1}) - K_2^{\top}(\pi_{2^m-2}) = K_2^{\perp}(\pi'_i) - K_2^{\top}(\pi'_{i+1}) \ , \end{aligned}$$

- for $i = 2^m - 1$:

$$\begin{aligned} & v_2(\pi'_i) - v_2(\pi'_{i+1}) = v_2(\pi_1) - v_2(\pi_{2^m-1}) \\ = \; & -x - (2^m - 3)\delta > -x - (2(2^m - 2))\delta \\ = \; & K_2^{\perp}(\pi_1) - K_2^{\top}(\pi_{2^m-1}) = K_2^{\perp}(\pi'_i) - K_2^{\top}(\pi'_{i+1}) \ . \end{aligned}$$

This ends the proof that $v_2 \in \mathsf{UD}_2(K_2)$ owing to Lemma C.3. □

D.2 Putting Things Together

Let the first two players respectively have candidate sets K_1 and K_2 and play the undominated strategies v_1 and v_2 (from Claim D.1 and Claim D.2, and see also Figure 2); let the rest of the players have valuation 0 and report 0 (which is an undominated strategy for each such player).

We make the following observations:

- When the players report $v \stackrel{\text{def}}{=} (v_1, v_2, 0, \dots, 0)$, the VCG mechanism will always choose the allocation $A = ([m], \varnothing, \dots, \varnothing)$.

  Indeed, the social welfare of A relative to v is

  $$v_1([m]) = v_1(\pi_{2^m-1}) = x + (2^m - 2)\delta \ .$$

  On the other hand, for any allocation giving $\pi_i \neq \varnothing$ to player 1 and $\pi_{2^m-1-i} = \overline{\pi_i}$ to player 2, the social welfare relative to v is equal to

  $$v_1(\pi_i) + v_2(\pi_{2^m-1-i}) = (x + (i-1)\delta) + (2^m - 1 - i)\delta - \varepsilon = x + (2^m - 2)\delta - \varepsilon \ ,$$

  which is smaller than that achieved by A; furthermore, for any allocation giving $\varnothing$ to player 1 and $[m]$ to player 2, the social welfare relative to v is equal to

  $$v_2([m]) = v_2(\pi_{2^m-1}) = x + (2^m - 2)\delta - \varepsilon \ ,$$

  which again is also smaller than that achieved by A.

- Assume that we pick the true valuation $\theta_1 \in K_1$ for player 1 to be such that $\theta_1(S) = x$ for all non-empty S, and $\theta_2 \in K_2$ for player 2 to be such that $\theta_2(S) = K_2^\top(S)$. Of course, we can only choose $\theta_i(S) = 0$ for all other players $i > 2$. (See Figure 2)
- The true social welfare on allocation A is $\theta_1([m]) = x$.
- The maximum social welfare is instead the following:

$$\mathrm{MSW}(\theta) \geq \theta_2([m]) = K_2^\top(\pi_{2^m-1}) = x + (2(2^m - 2) + 1)\delta - \varepsilon \ .$$

- Hence, the obtained social welfare compared to the maximum social welfare in this case is

$$\mathrm{SW}(\theta, \mathsf{VCG}(v)) = x \leq \mathrm{MSW}(\theta) - (2(2^m - 2) + 1)\delta + \varepsilon \ .$$

  By choosing $\varepsilon > 0$ sufficiently small, the social welfare guarantee of the VCG mechanism is at most

$$\mathrm{MSW}(\theta) - (2^{m+1} - 3)\delta \ .$$

This finishes the proof of (D.2), the worst-case choice of θ for Theorem 1b.

For the best-case choice of θ, we observe that for the same choice of v_1, v_2, K_1, K_2, A:

- The true social welfare on allocation A is $\theta_1([m]) \leq x + \delta/2$.
- The maximum social welfare is instead the following:

$$\mathrm{MSW}(\theta) \geq \theta_2([m]) \geq K_2^\perp(\pi_{2^m-1}) = x + 2(2^m - 2)\delta - \varepsilon \ .$$

- Hence, the obtained social welfare compared to the maximum social welfare in this case is

$$\mathrm{SW}(\theta, \mathsf{VCG}(v)) \leq x + \delta/2 \leq \mathrm{MSW}(\theta) - 2(2^m - 2)\delta + \delta/2 + \varepsilon \ .$$

  By choosing $\varepsilon > 0$ sufficiently small, the social welfare guarantee of the VCG mechanism is at most

$$\mathrm{MSW}(\theta) - (2^{m+1} - 5)\delta \ .$$

This finishes the proof of (D.1), the best-case choice of θ for Theorem 1b. ■

E Theorem 2 with Mixed Strategies

In this section we prove an analogue of Theorem 2 for mixed strategies, as follows.

Theorem 2′. *In a combinatorial Knightian auction with n players and m goods, let the VCG mechanism break ties by preferring subsets with smaller cardinalities.*[28] *Then, for all δ, all products K of δ-approximate candidate sets, all profiles $\theta \in K$, all profiles of mixed strategies $\sigma \in \mathsf{RM}^{\mathsf{mix}}(K)$, and all $p \geq 1$, we have with probability at least $1 - 1/p$ over the choices of v from σ:*

$$\mathrm{SW}(\theta, \mathsf{VCG}(v)) \geq \mathrm{MSW}(\theta) - O(n^2 p) \cdot \delta \ .$$

(This result can be tightened to $O(n \log n \log(1/p) \cdot \delta)$ either when (1) players are restricted to consider only monotone valuations (i.e., $\theta_i(S) \leq \theta_i(T)$ for any $S \subseteq T$), or when (2) players are studying $\mathsf{RM}^{\mathsf{mix}}(\mathsf{UD}(K))$ strategies, rather than just $\mathsf{RM}^{\mathsf{mix}}(K)$.)

Before proving this theorem, we first illustrate why the result is very different from that of Theorem 2.

E.1 Why Allowing Mixed Strategies Yields a Different Result

When a regret-minimizing player considers mixed strategies, he may significantly deviate (in expectation) from his candidate set. (This stands in contrast to the pure-strategy case, where he may deviate by at most δ; cf. Claim 5.2.) In fact, deviating may happen even in a single-good auction.

An Example in a Single-Good Auction. Let i be a player with candidate set $K_i = [x, x+\delta]$ in a single-good (Knightian) auction. One can carefully verify that his minimum regret is at most $\frac{\delta}{4}$, obtained by a mixed strategy of bidding uniformly at random between x and $x+\delta$. However, we state without proof that the following mixed strategy σ_i also provides a regret of $\frac{\delta}{4}$:

$$\sigma_i = \begin{cases} \text{drawn uniformly at random from } [x, x+\frac{3}{4}\delta] & \text{w.p. } \frac{3}{4}; \\ x + t\delta, & \text{w.p. } \frac{1}{4}(\frac{1}{t} - \frac{1}{t+1}) \text{ where } t \in \mathbb{Z}_+. \end{cases} \tag{E.1}$$

Note that the expected bidding value $\mathbb{E}[\sigma_i] = +\infty$ is unbounded from above, and one can similarly construct a strategy in which player i arbitrarily (in expectation) underbids. This destroys the hope of using linearity of expectation to deduce the mixed-strategy case as a corollary of the pure-strategy one.

However, any such deviation always satisfies the probabilistic guarantee $\Pr[\sigma_i \geq x + t\delta] \leq \frac{1}{4t}$ for overbidding (and similarly, underbidding), resulting in the simple conclusion that, with constant probability, none of the n players over/underbids by more than $O(n\delta)$. The social welfare is therefore affected by at most $O(n^2\delta)$ in a single-good auction.[29]

A Harder Problem in Combinatorial Auctions. In combinatorial auctions with m goods, each player reports $2^m - 1$ values on each of the $2^m - 1$ non-empty subsets of $[m]$. Thus, a player may (in principle) choose to independently overbid or underbid each of his $2^m - 1$ coordinates, according to (E.1). If so, then, with constant probability, he may choose to (a) overbid by $O(2^m\delta)$ on one of his coordinates, and (b) underbid by $O(2^m\delta)$ on another.

This possibility complicates the analysis, because such a choice of strategy may lead to a social welfare loss of $O(2^m\delta)$. Interestingly, we show that (a) cannot happen, but (b) can. However, when (b) happens, the social welfare is not going to be affected much.

[28] If giving subsets A or $B \subsetneq A$ to player i provides the same social welfare, then the VCG will give B to player i.

[29] A more careful analysis leads to $O(n \log n \cdot \delta)$.

E.2 Proof of Theorem 2′

Proof. We begin by explicitly writing down the formulation of the (maximum) regret in (5.1) for mixed strategies. Given a candidate set K_i of player i, and a possibly mixed strategy σ_i from which his bidding strategy v_i is drawn, the (expected maximum) regret of σ_i for player i is

$$R_i(K_i, \sigma_i) = \max_{\theta_i \in K_i} \max_{v_{-i}} \Big(\mathrm{MSW}\big(\theta_i, v_{-i}\big) - \mathbb{E}_{v_i \sim \sigma_i}\big[\mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big)\big] \Big) \ . \tag{E.2}$$

We also recall the following notations. For each player i, each candidate set $K_i \subset \Theta_i$, and each subset $T \subseteq [m]$, we let

$$K_i(T) \stackrel{\text{def}}{=} \{\theta_i(T)\}_{\theta_i \in K_i}, \quad K_i^{\perp}(T) \stackrel{\text{def}}{=} \inf K_i(T),$$
$$K_i^{\top}(T) \stackrel{\text{def}}{=} \sup K_i(T), \quad K_i^{\text{mid}}(T) \stackrel{\text{def}}{=} (K_i^{\perp}(T) + K_i^{\top}(T))/2 \ .$$

For the same reason as Footnote 16 on page 13 in the main paper, we assume without loss of generality that for each T, the minimum/maximum point in $K_i(T)$ exists. That is, $K_i^{\perp}(T) \stackrel{\text{def}}{=} \min K_i(T)$ and $K_i^{\top}(T) \stackrel{\text{def}}{=} \max K_i(T)$.

We first note that Claim 5.1 continues to hold:[30]

Claim 5.1. *Let v_i be a strategy of player i such that $v_i(T) = K_i^{\text{mid}}(T)$ for each non-empty $T \subseteq [m]$. Then $R_i(K_i, v_i) \le \delta$.*

We now prove some properties about an arbitrary (possibly mixed) strategy σ_i of player i with regret $\le \delta$.

Player Underbidding

We first show a variant of Claim 5.2a from the main paper. It is a probabilistic bound on how a player i may underbid on each of his $2^m - 1$ coordinates:

Claim E.1 (player underbidding)**.** *Let σ_i be a (possibly mixed) strategy of player i such that $R_i(K_i, \sigma_i) \le \delta$. Then, for any non-empty subset $T \subseteq [m]$, and any real number $t \ge 1$,*

$$\Pr_{v_i \sim \sigma_i}\left[K_i^{\top}(T) - \max_{T' \subseteq T} v_i(T') > t \cdot \delta\right] \le \frac{1}{t} \ .$$

Proof. Suppose the claim is not true. Then, there exists T such that

$$\Pr_{v_i \sim \sigma_i}\left[K_i^{\top}(T) - \max_{T' \subseteq T} v_i(T') > t \cdot \delta\right] > \frac{1}{t} \ . \tag{E.3}$$

We contradict our assumption on v_i by showing $R_i(K_i, \sigma_i) > \delta$.

To show $R_i(K_i, \sigma_i) > \delta$, as per (E.2), we must find some v_{-i} and some θ_i so that

$$\mathrm{MSW}\big(\theta_i, v_{-i}\big) - \mathbb{E}_{v_i \sim \sigma_i}\big[\mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big)\big] > \delta \tag{E.4}$$

[30] We note that when mixed strategies are allowed, one can find a strategy with regret $\delta/2$, therefore bidding the mid-points, having a regret δ, is no longer a regret-minimizing strategy. Since the remaining proof of Theorem 2′ only requires to know that 'the regret-minimizing strategy has a regret $O(\delta)$', it suffices to analyze the mid-points, losing a constant factor of 2.

Let j be an arbitrary player other than i. We choose $\theta_i \in K_i$ such that $\theta_i(T) = K_i^\top(T)$ and v_{-i} as follows: for every $S \subseteq [m]$

$$v_j(S) \stackrel{\text{def}}{=} \begin{cases} H & \text{if } S = \overline{T} \\ H + (K_i^\top(T) - t \cdot \delta) & \text{if } S = [m] \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad v_k(S) \stackrel{\text{def}}{=} 0 \text{ for every } k \notin \{i,j\}.$$

Above, H is some huge real number (i.e., much bigger than $v_i(S)$ for any subset S).[31]

Recall that (E.3) tells us that, with probability more than $\frac{1}{t}$ over the choice of v_i from σ_i, the event $K_i^\top(T) - \max_{T' \subseteq T} v_i(T') > t \cdot \delta$ occurs. Let us denote by $\textsc{Event}(v_i)$ this event, and it is not hard to verify that $\textsc{Event}(v_i)$ implies that the outcome $\mathsf{VCG}(v_i, v_{-i})$ must allocate $\varnothing$ to player i, and $[m]$ to player j. Therefore, with probability more than $\frac{1}{t}$, we have

$$\mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) = \theta_i(\varnothing) + v_{-i}([m]) = H + K_i^\top(T) - t \cdot \delta \ .$$

On the other hand, $\mathrm{MSW}(\theta_i, v_{-i}) \geq \theta_i(T) + v_{-i}(\overline{T}) = K_i^\top(T) + H$, and therefore

$$\begin{aligned} &\mathbb{E}_{v_i \sim \sigma_i}\big[\mathrm{MSW}\big(\theta_i, v_{-i}\big) - \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big)\big] \\ &\geq \Pr_{v_i \sim \sigma_i}[\textsc{Event}(v_i)] \cdot \mathbb{E}_{v_i \sim \sigma_i}\Big[\mathrm{MSW}(\theta_i, v_{-i}) - \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) \Big| \textsc{Event}(v_i)\Big] \\ &> \frac{1}{t} \cdot \Big(K_i^\top(T) + H - \big(H + (K_i^\top(T) - t \cdot \delta)\big)\Big) = \delta \ . \end{aligned}$$

This proves (E.4) and concludes concludes the proof of Claim E.1. □

We remark here that the above proof matches our high level description in Appendix E.1. That is, since a player may have different valuations on all of his $2^m - 1$ coordinates, he may choose to independently underbid each of his $2^m - 1$ coordinates according to Claim E.2 (which is tight, due to an example generalizing (E.1) to allow multiple goods). If so, with constant probability (using union bound), he may underbid by $O(2^m \delta)$ on one of his $2^m - 1$ coordinates.

Could this large underbidding destroy the social welfare by $O(2^m \delta)$? Our answer is No (as we shall formally explain later) because, if, in the maximum social welfare allocation, player i receives a subset $B_i \subseteq [m]$ of the goods, all we need to learn from the player's underbidding is: *how much will player i underbid on coordinate B_i?* Therefore, we do not care how much he underbids on other coordinates, and therefore this 2^m factor does not show up in the social welfare loss.

Player Overbidding

The overbidding case is much harder. In fact, one can (essentially) show a similar coordinate-wise argument as in Claim E.1, and conclude that a player will overbid on each of his coordinates by at most $t \cdot \delta$, with probability at most $\frac{1}{t}$. Via a union bound, this implies that, with constant probability, he may overbid by $O(2^m \delta)$ on *one* of his $2^m - 1$ coordinates. If this happens, unlike the underbidding case, the social welfare performance will be very poor. The following example illustrates this point.

> Example. Consider a 2-player auction with m goods, where m is even. The first player is only interested in the subsets of $[m]$ that have cardinality $m/2$, and his value for each such subset lies in the interval $[x, x+\delta]$. The second player is only interested in the set of all goods,

[31] Notice that when $T = [m]$ we have $\overline{T} = \varnothing$, and one cannot assign $v_j(\varnothing)$ to be a nonzero number. In that case, we can choose $H = 0$ and v_j remains well-defined, since we must have $K_i^\top(T) - t \cdot \delta > 0$ (as otherwise $K_i^\top(T) - \max_{T' \subseteq T} v_i(T') > t \cdot \delta$ cannot hold, contradicting our assumption). The rest of the proof still goes through.

$[m]$, which he values precisely $x + (\binom{m}{m/2} - 1)\delta$. Notice that the maximum social welfare in this setting is $x + (\binom{m}{m/2} - 1)\delta$. Also notice that, in such an auction, at most one player 'wins'. That is, at most one player can be allocated a subset of $[m]$ which he positively values.

Now suppose that player 2 reports his true valuation, while player 1 overbids as follows. Let $t = \binom{m}{m/2}$. For each of the t subsets he is interested in, player 1 reports, independently and with probability $1/t$, the value $x + t \cdot \delta$, and x otherwise. (For each subset he is not interested in, player 1 reports 0.) Then, with constant probability, player 1 reports $x + t \cdot \delta$ on one of his coordinates, and thus 'wins' the auction. Note that, when player 1 'wins', the social welfare is at most $x + \delta$ and misses the maximum social welfare by $(t-2) \cdot \delta = \tilde{\Omega}(2^m \delta)$.

Therefore, to prove a good social-welfare performance, it is not advisable to bound a player's overbidding *coordinate-wise*. In fact, we prove the following claim, which is significantly different from what we showed in Claim 5.2b for the pure case. The new claim essentially bounds how a player i may overbid (on all coordinates) with respect to a given mixed strategy sub-profile σ_{-i} of his opponents. Since we will eventually be interested in only *one* particular σ_{-i} —namely, the one when all players other than i are playing regret-minimizing strategies— we do not need to pay for the extra $O(2^m \delta)$ loss in the union bound.

Claim E.2 (player overbidding)**.** *Let σ_i be a (possibly mixed) strategy of player i such that $R_i(K_i, \sigma_i) \le \delta$, σ_{-i} an arbitrary (possibly mixed) strategy sub-profile of his opponents, and $\theta_i \in K_i$ his possible true valuation. Then, for any real number $t \ge 1$,*

$$\Pr_{\substack{v_i \sim \sigma_i \\ v_{-i} \sim \sigma_{-i}}} \Big[v_i\big(\mathsf{VCG}(v_i, v_{-i})\big) > \theta_i\big(\mathsf{VCG}(v_i, v_{-i})\big) + 4t \cdot \delta \Big] \le \frac{1}{t} \ . \tag{E.5}$$

Proof. Suppose the claim is not true and there are choices of σ_i, σ_{-i}, and θ_i, such that the above probability is strictly larger than $\frac{1}{t}$. We denote by $\textsc{Event}_1(v_i, v_{-i})$ the probabilistic event that $v_i\big(\mathsf{VCG}(v_i, v_{-i})\big) > \theta_i\big(\mathsf{VCG}(v_i, v_{-i})\big) + 4t \cdot \delta$, and we want to show that if $\Pr[\textsc{Event}_1] > \frac{1}{t}$, then $R_i(K_i, \sigma_i) > \delta$, contradicting our assumption on σ_i. To achieve this, we lower bound (E.2) (using the same choice of θ_i provided in the assumption of this claim) by a probabilistic form:

$$R_i(K_i, \sigma_i) \ge \mathbb{E}_{v^*_{-i} \sim \sigma^*_{-i}} \Big[\mathrm{MSW}\big(\theta_i, v^*_{-i}\big) - \mathbb{E}_{v_i \sim \sigma_i} \big[\mathrm{SW}\big((\theta_i, v^*_{-i}), \mathsf{VCG}(v_i, v^*_{-i})\big)\big] \Big] \ . \tag{E.6}$$

Now it suffices to choose a witness distribution σ^*_{-i} so that the right-hand side is larger than δ.

We choose σ^*_{-i} as follows. It is reconstructed from the distribution σ_{-i} given in the assumption, with every occurrence of $v_{-i} \sim \sigma_{-i}$ replaced by v^*_{-i} with the same probability, where v^*_{-i} is defined as:

$$\forall j \ne i \, \forall S \subseteq [m] \quad v^*_j(S) \stackrel{\text{def}}{=} \begin{cases} \mathrm{MSW}(\theta_i, v_{-i}) + 2t \cdot \delta & \text{if } S = [m] \\ v_j(S) & \text{otherwise} \end{cases} \ .$$

Now assuming, by way of contradiction, that the desired regret term $R_i(K_i, \sigma_i) \le \delta$, which implies (using (E.2) for v_{-i} drawn from σ_{-i}):

$$\mathbb{E}_{v_{-i} \sim \sigma_{-i}} \Big[\mathrm{MSW}\big(\theta_i, v_{-i}\big) - \mathbb{E}_{v_i \sim \sigma_i} \big[\mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big)\big] \Big] \le R_i(K_i, \sigma_i) \le \delta \ .$$

Using Markov bound, with probability at least $1/2t$ over the choices of $v_i \sim \sigma_i$ and $v_{-i} \sim \sigma_{-i}$, we have

$$\mathrm{MSW}\big(\theta_i, v_{-i}\big) - \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) \le 2t \cdot \delta$$

We denote by $\textsc{Event}_2(v_i, v_{-i})$ the probabilistic event such that the above inequality is true. From (E.5), we know that with probability strictly larger than $1/t - 1/2t = 1/2t$ we have that both $\textsc{Event}_1$ and $\textsc{Event}_2$ happen, and therefore

$$
\begin{aligned}
& v_i\big(\mathsf{VCG}(v_i, v_{-i})\big) > \theta_i\big(\mathsf{VCG}(v_i, v_{-i})\big) + 4t \cdot \delta && \text{(using } \textsc{Event}_1\text{)} \\
\implies & v_i\big(\mathsf{VCG}(v_i, v_{-i})\big) + v_{-i}\big(\mathsf{VCG}(v_i, v_{-i})\big) > \theta_i\big(\mathsf{VCG}(v_i, v_{-i})\big) + v_{-i}\big(\mathsf{VCG}(v_i, v_{-i})\big) + 4t \cdot \delta && \\
\implies & \mathrm{SW}\big((v_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i}\big) > \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) + 4t \cdot \delta && \\
\implies & \mathrm{MSW}(v_i, v_{-i}) > \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) + 4t \cdot \delta && \\
\implies & \mathrm{MSW}(v_i, v_{-i}) > \mathrm{MSW}(\theta_i, v_{-i}) + 2t \cdot \delta && \text{(using } \textsc{Event}_2\text{)} \\
\implies & \mathrm{MSW}(v_i, v_{-i}) > v_j^*([m]) && (\forall j \neq i, \text{ using the definition of } v_{-i}^* \text{ .)}
\end{aligned}
$$

The last strict inequality implies that the allocation under $\mathsf{VCG}(v_i, v_{-i})$ must be the same as $\mathsf{VCG}(v_i, v_{-i}^*)$. This is because v_{-i}^* is only different from v_{-i} on the coordinates $[m]$ for players $j \neq i$, but those coordinates only incur a smaller social welfare than $\mathsf{VCG}(v_i, v_{-i})$ according to the last inequality above.

In sum, we have that $\mathrm{MSW}(\theta_i, v_{-i}^*) \geq v_{-i}([m]) = \mathrm{MSW}(\theta_i, v_{-i}) + 2t \cdot \delta$; however, under $\textsc{Event}_1 \wedge \textsc{Event}_2$, the obtained social welfare can be upper bounded as follows:

$$
\begin{aligned}
\mathrm{SW}\big((\theta_i, v_{-i}^*), \mathsf{VCG}(v_i, v_{-i}^*)\big) &= \mathrm{SW}\big((\theta_i, v_{-i}^*), \mathsf{VCG}(v_i, v_{-i})\big) = \mathrm{SW}\big((\theta_i, v_{-i}), \mathsf{VCG}(v_i, v_{-i})\big) \\
&\leq \mathrm{MSW}(\theta_i, v_{-i}) \leq \mathrm{MSW}(\theta_i, v_{-i}^*) - 2t \cdot \delta \ .
\end{aligned}
$$

Above, the first equality is because $\mathsf{VCG}(v_i, v_{-i})$ produces the same allocation as $\mathsf{VCG}(v_i, v_{-i}^*)$; the second equality is because $\mathsf{VCG}(v_i, v_{-i}^*)$ never gives all the goods to a player $j \neq i$; and the first inequality is because, by definition, the VCG maximizes social welfare.

Now we go back to (E.6), and show that $R_i(K_i, \sigma_i) > \delta$:

$$
\begin{aligned}
R_i(K_i, \sigma_i) &\geq \mathbb{E}_{\substack{v_i \sim \sigma_i \\ v_{-i}^* \sim \sigma_{-i}^*}} \Big[\mathrm{MSW}\big(\theta_i, v_{-i}^*\big) - \mathrm{SW}\big((\theta_i, v_{-i}^*), \mathsf{VCG}(v_i, v_{-i}^*)\big)\Big] \\
&\geq \Pr_{\substack{v_i \sim \sigma_i \\ v_{-i}^* \sim \sigma_{-i}^*}} [\textsc{Event}_1 \wedge \textsc{Event}_2] \times \\
&\quad \mathbb{E}_{\substack{v_i \sim \sigma_i \\ v_{-i}^* \sim \sigma_{-i}^*}} \Big[\mathrm{MSW}\big(\theta_i, v_{-i}^*\big) - \mathrm{SW}\big((\theta_i, v_{-i}^*), \mathsf{VCG}(v_i, v_{-i}^*)\big) \Big| \textsc{Event}_1 \wedge \textsc{Event}_2\Big] \\
&> \frac{1}{2t} \times 2t \cdot \delta = \delta \ .
\end{aligned}
$$

The above conclusion contradicts our assumption that the regret of the mixed strategy σ_i is at most δ. This concludes the proof of the claim. □

Putting It All Together

Now we go back to the proof of Theorem 2. Let $\sigma = (\sigma_1, \ldots, \sigma_n) \in \mathsf{RM}^{\mathsf{mix}}(K)$ be a profile of regret-minimizing mixed strategies, and let $\theta \in K$ be any valuation profile. Since there exists a strategy with regret $\leq \delta$ for each player (see Claim 5.1), we must have $R_i(K_i, \sigma_i) \leq \delta$ to satisfy the assumption of Claim E.1 and E.2.

Now, letting $(B_0, B_1, \ldots, B_n)$ be the allocation that maximizes the social welfare under θ, we are ready to compute the social welfare guarantee. For any choice of $v \sim \sigma$, let X_i denote the non-negative probabilistic variable equal to the difference $v_i\big(\mathsf{VCG}(v)\big) - \theta_i\big(\mathsf{VCG}(v)\big)$; according to Claim E.2, we have $\Pr[X_i > 4t\delta] < \frac{1}{t}$. Also let Y_i denote the non-negative probabilistic variable equal to the difference $K_i^\top(B_i) - \max_{T' \subseteq B_i} v_i(T')$, and, according to Claim E.1 (for the choice of $T = B_i$), we have $\Pr[Y_i > t\delta] \le \frac{1}{t}$.

$$\begin{aligned} \mathrm{SW}(\theta, \mathsf{VCG}(v)) &= \sum_{i=1}^n \theta_i(\mathsf{VCG}(v)) = \sum_{i=1}^n v_i(\mathsf{VCG}(v)) - \sum_{i=1}^n X_i \\ &\ge \sum_{i=1}^n \max_{T' \subseteq B_i} v_i(T') - \sum_{i=1}^n X_i \qquad \text{(because the VCG maximizes social welfare under } v\text{)} \\ &= \sum_{i=1}^n K_i^\top(B_i) - \sum_{i=1}^n (X_i + Y_i) \\ &\ge \sum_{i=1}^n \theta_i(B_i) - \sum_{i=1}^n (X_i + Y_i) = \mathrm{MSW}(\theta) - \sum_{i=1}^n (X_i + Y_i) \ . \end{aligned}$$

We are now left to bound $\sum_{i=1}^n (X_i + Y_i)$. For any $p \ge 1$ and each choice of $i \in [n]$, with probability at least $1 - \frac{1}{2np}$, we have that $X_i \le (8np)\delta$, and, with probability at least $1 - \frac{1}{2np}$, $Y_i \le (2np)\delta$. Using union bound, with a total probability of at least $1 - \frac{1}{p}$ (over the choices of v from σ), we have $X_i \le (8np)\delta$ and $Y_i \le (2np)\delta$ *for all* $i \in [n]$. In such a case the above difference satisfies

$$\mathrm{SW}(\theta, \mathsf{VCG}(v)) \ge \mathrm{MSW}(\theta) - O(n^2 p) \cdot \delta \ .$$

This concludes the proof of Theorem 2′. ■

References

[AM92] Dilip Abreu and Hitoshi Matsushima. Virtual implementation in iteratively undominated strategies: Complete information. *Econometrica*, 60(5):993–1008, 1992.

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

[BDF$^+$10] Dave Buchfuhrer, Shaddin Dughmi, Hu Fu, Robert Kleinberg, Elchanan Mossel, Christos Papadimitriou, Michael Schapira, Yaron Singer, and Chris Umans. Inapproximability for vcg-based combinatorial auctions. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 518–536, 2010.

[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.

[BLP06] Moshe Babaioff, Ron Lavi, and Elan Pavlov. Single-value combinatorial auctions and implementation in undominated strategies. In *SODA '06: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1054–1063, New York, NY, USA, 2006. ACM.

[Bod12] Aaron L. Bodoh-Creed. Ambiguous beliefs and mechanism design. *Games and Economic Behavior*, 75(2):518–537, 2012.

[BOP06] Subir Bose, Emre Ozdenoren, and Andreas Pape. Optimal auctions with ambiguity. *Theoretical Economics*, 1(4):411–438, December 2006.

[CAT09] C. G. Chorus, T. A. Arentze, and H. J. P. Timmermans. Spatial choice: a matter of utility or regret? *Environment and Planning B: Planning and Design*, 36(3):538–551, 2009.

[CMZ12] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Mechanism design with approximate valuations. In *Proceedings of the 3rd Innovations in Theoretical Computer Science conference*, ITCS '12, 2012.

[CMZ14a] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Bridging utility maximization and regret minimization. *ArXiv e-prints*, abs/1403.6394, March 2014. `http://arxiv.org/abs/1403.6394`.

[CMZ14b] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Knightian analysis of the VCG mechanism in unrestricted combinatorial auctions. *ArXiv e-prints*, abs/1403.6410, March 2014. `http://arxiv.org/abs/1403.6410`.

[CMZ14c] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Knightian robustness from regret minimization. *ArXiv e-prints*, abs/1403.6409, March 2014. `http://arxiv.org/abs/1403.6409`.

[CMZ14d] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Knightian robustness of single-parameter domains. *ArXiv e-prints*, abs/1403.6411, March 2014. `http://arxiv.org/abs/1403.6411`.

[CMZ14e] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Knightian robustness of the Vickrey mechanism. *ArXiv e-prints*, abs/1403.6413, March 2014. `http://arxiv.org/abs/1403.6413`.

[CS05] Vincent Conitzer and Tuomas Sandholm. Complexity of (iterated) dominance. In *Proceedings of the 6th ACM conference on Electronic commerce*, EC '05, pages 88–97, New York, NY, USA, 2005. ACM.

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

[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.

[DV11] Shaddin Dughmi and Jan Vondrák. Limitations of randomized mechanisms for combinatorial auctions. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 502–511. IEEE, 2011.

[EK07] Richard Engelbrecht-Wiggans and Elena Katok. Regret in auctions: Theory and evidence. *Economic Theory*, 33(1):81–101, 2007.

[Eng89] Richard Engelbrecht-Wiggans. The effect of regret on optimal bidding in auctions. *Management Science*, 35(6):685–692, 1989.

[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.

[FO07] Emel Filiz and Erkut Y Ozbay. Auctions with anticipated regret: Theory and experiment. *The American Economic Review*, 97(4):1407–1418, 2007.

[FT91] Drew Fudenberg and Jean Tirole. *Game theory. 1991*. MIT Press, 1991.

[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.

[GM75] 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.

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

[HB04] Nathanael Hyafil and Craig Boutilier. Regret Minimizing Equilibria and Mechanisms for Games with Strict Type Uncertainty. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 268–277, July 2004.

[HGC11] David A Hensher, William H Greene, and Caspar G Chorus. Random regret minimization or random utility maximization: an exploratory analysis in the context of automobile fuel choice. *Journal of Advanced Transportation*, 2011.

[HP12] Joseph Y. Halpern and Rafael Pass. Iterated regret minimization: A new solution concept. *Games and Economic Behavior*, 74(1):184–207, January 2012. A preliminary version appeared in IJCAI'09.

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

[JPS94] Matthew O. Jackson, Thomas Palfrey, and Sanjay Srivastava. Undominated Nash implementation in bounded mechanisms. *Games and Economic Behavior*, 6(3):474–501, 1994.

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

[LR89] Peter B Linhart and Roy Radner. Minimax-regret strategies for bargaining over several variables. *Journal of Economic Theory*, 48(1):152–178, 1989.

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

[LS08] Kevin Leyton-Brown and Yoav Shoham. Essentials of game theory: A concise multidisciplinary introduction. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2(1):1–88, 2008.

[Mas74] Andrew Mas-Colell. An equilibrium existence theorem without complete or transitive preferences. *Journal of Mathematical Economics*, 1(3):237–246, December 1974.

[Mil54] John W. Milnor. Games against nature. In Robert M Thrall, Clyde Hamilton Coombs, and Robert L Davis, editors, *Decision processes*. John Wiley & Sons, Inc., 1954.

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

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

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

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

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

[RS10] Ludovic Renou and Karl H. Schlag. Minimax regret and strategic uncertainty. *Journal of Economic Theory*, 145(1):264–286, January 2010.

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

[Sav51] Leonard J Savage. The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67, 1951.

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

[Sel89] Reinhard Selten. Blame avoidance as motivating force in the first price sealed bid private value auction. In *Economics Essays in Honor of Werner Hildenbrand*, pages 333–344. Springer, 1989.

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

[Sto11] Joerg Stoye. Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146(6):2226–2251, 2011.

[TLB07] David R. M. Thompson and Kevin Leyton-Brown. Valuation uncertainty and imperfect introspection in second-price auctions. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 1*, AAAI '07, pages 148–153, 2007.

[Wal49] Abraham Wald. Statistical decision functions. *The Annals of Mathematical Statistics*, 20(2):pp. 165–205, 1949.