

Mechanism Design with Approximate Types

by

Zeyuan Allen Zhu

B.S. in Mathematics and Physics, Tsinghua University (2010)

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering and Computer Science


at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2012

© Zeyuan Allen Zhu, MMXII. All rights reserved.

ARCHIVES



The author hereby grants to MIT permission to reproduce and
distribute publicly paper and electronic copies of this thesis document
in whole or in part.

A handwritten signature in black ink, appearing to read "Zeyuan Allen Zhu".

Author
Department of Electrical Engineering and Computer Science
January 6, 2012

Certified by
Silvio Micali
Professor
Thesis Supervisor

Accepted by
Leslie A. Kolodziejcki
Chairman, Department Committee on Graduate Theses

Mechanism Design with Approximate Types

by

Zeyuan Allen Zhu

Submitted to the Department of Electrical Engineering and Computer Science
on January 6, 2012, in partial fulfillment of the
requirements for the degree of
Master of Science in Electrical Engineering and Computer Science

Abstract

In mechanism design, we replace the strong assumption that each player knows his own payoff type *exactly* with the more realistic assumption that he knows it only *approximately*: each player i only knows that his true type θ_i is one among a set K_i , and adversarially and secretly chosen in K_i at the beginning of the game.

This model is closely related to the Knightian [20] notion of uncertainty in economics, but we consider it from purely mechanism design's perspective. In particular, we study the classical problem of maximizing social welfare in auctions when players know their true valuations only within a constant multiplicative factor $\delta \in (0, 1)$.

For single good auctions, we prove that no dominant-strategy mechanism can guarantee better social welfare than assigning the good to a random player. On the positive side, we provide tight upper and lower bounds for the social welfare achievable in undominated strategies, whether deterministically or probabilistically.

For multiple-good auctions, we prove that all dominant-strategy mechanisms can guarantee only an exponentially small fraction of the maximum social welfare, and the celebrated VCG mechanism (which is no longer dominant-strategy) guarantees, in undominated strategies, at most a doubly exponentially small fraction.

For general games beyond auctions, we provide definitional foundations for this new approximate-type model, and provide a universality result showing that all reasonable (including Bayesian or Knightian) models of type uncertainty are equivalent to our set-theoretic one, at least for the setting when the type space is "convex".

This work was done in collaboration with Silvio Micali and Alessandro Chiesa.

Thesis Supervisor: Silvio Micali

Title: Professor

To my beloved mother, Xiaoli Xu

Acknowledgments

Tracing back to October 2009, when I first met Silvio Micali at Tsinghua University, I was deeply attracted by his inspiring and idiosyncratic talk in mechanism design, and entered this field. A year later in September 2010, I was luckily selected as a student of this adventurous and knowledgeable professor, and joined the CSAIL Theory Group. Since then, the work in this thesis has started and been under Silvio's careful and close supervision. I appreciate it more than I can say.

I would also like to thank Alessandro Chiesa who introduced me all secrets about MIT, so that I can integrate into this big family without difficulty. He is the one I spent the past 1.5 years discussing everything related or not related to this thesis with, the one that I learned all advanced techniques of \LaTeX from, the one that I complained the locomotive whistles near our office with, etc. Such collaboration is more enjoyable and engaging than ever.

Beyond this thesis, I am fortunate to have collaborated extensively with Jonathan Kelner, Martin Rinard, Sasa Misailovic and Christian Sommer, and more or less with many others at MIT. I am also grateful for the support that was not research-related, including that from our instructors, our students and our secretaries in the CSAIL Theory Group; and also including that from my sister-like friend Xiaodan Jia, who made me porridge and bought me balloons when I was extremely sick rushing before two paper deadlines.

This thesis is fully supported by a 9-month Greater China Fellowship, a 3-month Big George Ventures Fund, and a 4.5-month teaching assistantship for 6.854 at MIT. Therefore, I would like to thank the Greater China Fund, Ray Sidney and David Karger respectively to make my study at MIT possible.

Last but not least, I never forget to thank my dearest mother for her unceasing support and unselfish dedication to my family, and my education, throughout the past 23 years and 204 days. For all of the above, and the unnamed, please allow me to express my deepest gratitude, now and always.

Earlier Publications of this Work

Much of the material in this thesis has appeared in earlier publications done by the author in joint work with Alessandro Chiesa and Silvio Micali. The result of single-good auctions has appeared in ITCS 2012 [9]. The result of multi-good auctions has appeared in part as an internal technical report [7]. The result of general games has appeared in part as an MIT technical report [8]. The universality result of approximate types has never published before.

Contents

1	Introduction	13
1.1	Single-Good Auctions	13
1.1.1	Our Model of Self Uncertainty	14
1.1.2	Our Results	14
1.1.3	Our Techniques	15
1.2	Multi-Good Auctions	16
1.3	General Games	17
1.4	Universality of Approximate Types	19
1.5	Related Work	20
2	Single-Good Auctions	23
2.1	The Model	23
2.1.1	The Auction Context	23
2.1.2	The Mechanism	26
2.1.3	Solutions Concepts	27
2.1.4	The Performance Measure	27
2.2	Our Results	28
2.2.1	The Inadequacy of Dominant-Strategy Mechanisms	29
2.2.2	The Power of Deterministic Undominated-Strategy Mechanisms	30
2.2.3	The Greater Power of Probabilistic Undominated-Strategy Mechanisms	32
3	Single-Good Auctions: Techniques of Independent Interest	35
3.1	The Undominated Intersection Lemma	35
3.1.1	Details	37
3.2	The Distinguishable Monotonicity Lemma	40
3.2.1	Details	41
4	Single-Good Auctions: Missing Proofs	47
4.1	Proof of Theorem 1	47

4.2	Proof of Theorem 2b	49
4.3	Proof of Theorem 3b	50
4.4	Proof of Theorem 2a	52
4.5	Proof of Theorem 3a	53
4.5.1	A Very Restricted Search	53
4.5.2	Our Allocation Function	55
4.5.3	Our Mechanism $M_{\text{opt}}^{(\delta)}$	63
5	Multi-Good Auctions	65
5.1	The Model	65
5.2	Our Results	66
5.2.1	A First Impossibility Result	66
5.2.2	A Second Impossibility Result	67
6	Multi-Good Auctions:	
	Missing Proofs	69
6.1	Proof for Fact 5	69
6.2	Proof for Theorem 5	72
6.3	Proof for Theorem 6	75
6.3.1	Proof Sketch	75
6.3.2	The Bird Lemma	77
6.3.3	Proof of the Bird Lemma	78
6.3.4	The Hard Instance	89
6.3.5	Main Proof	92
7	General Games	95
7.1	Game: Context and Mechanism	95
7.2	Dominance	97
7.3	Solution Concepts	98
7.4	Implementation	100
7.5	New Revelation Principle	101
8	Universality of Approximate Types	103
9	Conclusion	111
A	Performance Diagrams	113
B	Our Optimal Mechanism $M_{\text{opt}}^{(\delta)}$	115

List of Figures

8-1 Three basic models of type uncertainty in our hierarchy. 107

A-1 Performance of our optimal mechanism 113

Chapter 1

Introduction

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

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

In this thesis, we explore this general topic from the following four major directions:

- Guaranteeing social welfare in auctions of a single good with n players.
- Guaranteeing social welfare in auctions of m goods with n players.
- Generalizing our notions to arbitrary games.
- Comparing our notions to other models of player uncertainty, such as Bayesian and Knightian.

1.1 Single-Good Auctions

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

Guaranteeing high social welfare in auctions of a single good with n players.

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

1.1.1 Our Model of Self Uncertainty

In a single-good auction, a player type is a natural number, referred to as a *valuation*. The possibility that a player in such an auction may not precisely know his own valuation strikes us to be quite realistic. No one would be too surprised if, tasked to figure out a firm’s *true* valuation for the good, different employees reported different values, or some of them reported ranges of values rather than single values. This said, for decision theory —let alone mechanism design!— to be meaningful at all, the players must know *something* about their own types. The classical work of [20] (see also [5]) envisages that each player knows that his own type is distributed according to one of several possible distributions. Other works —e.g., [14, 24, 30]— envisage a more structured kind of “self knowledge”: namely, each player knows *the* distribution from which his own type has been drawn.

Our model of the players’ “self uncertainty” is purely *set-theoretic*. In essence, each player i only knows that his own valuation θ_i is one of several candidates in a set K_i . In this model we investigate how good mechanisms one can design when each K_i is *relatively clustered*: that is, when “each player knows his own valuation within the same fixed percentage δ .”

1.1.2 Our Results

Our main results contain the following three theorems:

Theorem 1 (informal). *When $\delta > 0$, the best dominant-strategy mechanism can guarantee only a $\approx \frac{1}{n}$ fraction of the maximum social welfare.*

Theorem 2 (informal). *When $\delta \geq 0$, the best deterministic undominated-strategy mechanism is the second-price mechanism, that guarantees $\approx \left(\frac{1-\delta}{1+\delta}\right)^2$ fraction of the maximum social welfare.*

Theorem 3 (informal). *When $\delta > 0$, the best probabilistic undominated-strategy mechanism $M_{\text{opt}}^{(\delta)}$ outperforms the second-price one, and guarantees $\approx \frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}$ fraction of the maximum social welfare.*

The picture emerging from auctions of a single good is quite “reassuring”: dominant strategies are no longer attractive, but then the classical second-price mechanism (although developed for the simpler, exact-type setting) continues to deliver lots of efficiency in undominated strategies. Only if one insisted on optimal worst-case performance, should he consider our more complex probabilistic mechanism $M_{\text{opt}}^{(\delta)}$ from Theorem 3 above.

And of course, since we are in a new model, we need to carefully define notions like *dominant* and *undominated strategies*, as well as what it means by guaranteeing certain fraction of the maximum social welfare, when players do not even exactly know their true types.

1.1.3 Our Techniques

Exploring a new direction requires developing a new set of tools. To guide the design and analysis of mechanisms in our new setting, we find it useful to establish two structural lemmas. Note that the first one generalizes to *any* approximate-type game settings, and the second one generalizes to any *single-parameter* approximate-type game settings.

The Undominated Intersection Lemma. To prove impossibility results for general mechanisms, we must deal with arbitrary sets of strategies. In our model, even the “natural” strategies are richer than before. For instance, a player i having set K_i might report a single number (e.g., a member in K_i), or a set of numbers (e.g., K_i itself). However, no matter what strategies a mechanism may grant, we prove that for any two nontrivially overlapping sets K_i and K'_i , the corresponding sets of undominated strategies for player i must have a pair of (mixed) strategies that are “arbitrarily close”.

The Distinguishable Monotonicity Lemma. To prove possibility results, we are instead happy to consider only special sets of strategies. In particular, we consider mechanisms that constrain the players to report individual bids, just like the second-price mechanism. We shall prove that, if the allocation function of the mech-

anism grows monotonically with players' reports, plus some mild requirement, the mechanism will provide a clean characterization of the set of *undominated strategies* associated to any set K_i of a player i : namely, the strategies between the minimum and the maximum value in K_i .

1.2 Multi-Good Auctions

Now we raise the bar and consider a second goal (also trivially achieved when the players know their own types exactly):

Guaranteeing social welfare in combinatorial auctions¹ of m good with n players.

In combinatorial auctions, player i 's type θ_i is a function from all possible subsets of the goods to non-negative numbers, referred to as a *valuation*. As before, we assume that each player knows only some set K_i where $\theta_i \in K_i$, while each K_i is *relatively clustered*: that is, “each player knows his own valuation for each subset of the goods within the same fixed percentage δ .”

Our first impossibility result states that our negative result in Theorem 1 generalizes, with a exponentially worse bound in terms of the number of goods m .

Theorem 4 (informal). *When $\delta > 0$, the best dominant-strategy mechanism for combinatorial auctions can guarantee only a $\approx \frac{1}{n2^{m-1}}$ fraction of the maximum social welfare.*

Now we turn to the VCG mechanism [10, 18, 33], the celebrated extension of the second-price mechanism in combinatorial auctions. Recall that it is not a dominant-strategy mechanism anymore, and guarantees a reasonable fraction of the maximum social welfare in single-good approximate-valuation world (see Theorem 2). Should the similar result also hold in the multi-good case, mechanism design with approximate types would offer a pleasant off-the-road walk. Unfortunately, our reassuring picture in the single-good case vaporizes:

¹Here the term “combinatorial auction” refer to multi-good auctions without any restriction about player’s valuations. It is used in contract to for instance single-minded auctions, where players can have positive valuations to only one subset of the goods.

Theorem 6 (informal). *For any $n \geq 2$ and $m \geq 2$, the VCG mechanism can only guarantee $\left(\frac{1-\delta}{1+\delta}\right)^{2^m-2}$ fraction of the maximum social welfare.*

Extensions of this negative result, or even positive results but also utilizing player's external knowledge will be discussed in Chapter 5.

1.3 General Games

Our approximate valuation K_i for player i can be similarly defined in *any general* game, as long as K_i contains a list of possible candidate *types* of player i , rather than a list of valuations. We provide a self-contained chapter in this thesis for the notions of approximate-type games, along with the corresponding dominant and undominated implementations. Notice that a similar set-theoretic model but for a set of Bayesian distributions has been studied as a part of the Knightian decision theory, see Section 1.5 for comparisons.

Traditional definitions. Recall that in the *exact-type* world, for player i and his two strategies s_i and t_i , there are three naturally-defined notions of dominance. Informally,²

- s_i *very-weakly dominates* t_i if for *all* strategies of other players, $s_i \geq t_i$;
- s_i *weakly dominates* t_i if $\left\{ \begin{array}{l} \text{for all strategies of other players, } s_i \geq t_i \text{ and} \\ \text{for some strategy of other players, } s_i > t_i \end{array} \right.$;
- s_i *strictly dominates* t_i if for *all* strategies of other players, $s_i > t_i$.

Also, a strategy is said to be *dominant* if it very-weakly dominates everything else; and undominated if it is not weakly dominated by any other strategy.

Extensions to approximate types. A very straightforward generalization to the approximate-type world, at least for *very-weak* and *strict* dominance, is to require that $s_i \geq t_i$ or $s_i > t_i$ to hold not only for all strategy sub-profiles of other players, but also for *all possible type* $\theta_i \in K_i$. Notice that, when there is only one single

²We remark here that the names of the three notions are defined inconsistently among literatures, for instance, the very-weak dominance is also referred to weak dominance some times.

player, our quantifier on strategy sub-profiles of other players is *vacuous*, and this special case was also studied for instance by [1, 5, 11, 12, 17, 21, 25, 27, 29, 31]. To the best of our knowledge, in the approximate-type world, the notion of dominance taking into account strategy sub-profiles of other players was not studied before.

The generalization for *weak* dominance is more subtle. Essentially, we want to capture the weakest condition for t_i to be discarded in favor of s_i . We say then that weak dominance holds if very-weak dominance holds and, moreover, there exist a “witness” strategy sub-profile *and also* a “witness” type $\theta_i \in K_i$ such that $s_i > t_i$ is strict.

Mechanism design. Even in our close-related set-theoretic Knightian settings (see Section 1.5 for comparisons), little attempt has been made towards mechanism design (with an exception coming from the work Lopomo et al. [21]), and thus also towards understanding solution concepts and implementations in this setting. We provide modeling results in this direction.

In settings of incomplete information, two naturally-defined solution concepts are, informally speaking, “to guarantee desired property when players play dominant/undominated strategies”.

Implementations in dominant strategy. We say that a mechanism *implements* some property Π in dominant strategy if Π is achieved when each player plays a *specific* dominant strategy. For example, a *dominant strategy truthful (DST)* mechanism belongs to this category. In our approximate-type setting, a DST mechanism will ask each player to report a *set* of types, instead of a single type; a mechanism is DST and implements some property Π , if “reporting the truth K_i ” is (very-weakly) dominant for each player i , and Π is achieved under this strategy profile. We have been able to extend the revelation principle to the setting of approximate types:

Lemma 7.13 (Revelation Principle, Informal). *In the approximate-type world, if a property can be achieved via any dominant-strategy mechanism, it can also be achieved via a dominant-strategy-truthful mechanism.*

In the setting of a single player, our notion of DST mechanism is also studied by

Lopomo et al. [21]. However, we emphasize here a difference. If players are uncertain about their types, how to check if a property—which takes players’ type θ as an input—is achieved? In our paper, we use the worst-case analysis from computer science, and define it to be achieved only when for any possible type $\theta \in K$, the property holds; while in [21], they consider that the property holds at *some* possible type $\theta \in K$.

Implementation in undominated strategies. Since DST mechanisms are not always desirable (also when players know their types exactly), we also define and study *implementations* that guarantee a social property to hold whenever each player chooses an *arbitrary* strategy that is *not (weakly) dominated*. Similar ideas have already been used successfully (but when players know their types exactly) for designing mechanisms in single-value multi-minded auctions [2].

1.4 Universality of Approximate Types

For any general game, it is unrealistic not only to assume that “each player i knows *exactly* his own payoff type θ_i ”, but also “each player i knows an *exact* Bayesian prior D_i from which his own payoff type θ_i is drawn”. In principle, there are more than one possible candidates to model a player’s uncertainty without such *exact* knowledge:

- *Approximate Type Uncertainty.* The internal knowledge of each player i is modeled as a *set* of types K_i called the player’s *approximate type*.
- *Knightian Type Uncertainty.* The internal knowledge of each player i is modeled as a *set* \mathcal{D}_i of Bayesian distributions that is promised to contain the true distribution D_i^* where his own payoff type θ_i is drawn. This model was first introduced in economics by Knight [20], and then formalized by Bewley [5].
- *Hierarchy Type Uncertainty.* The internal knowledge of each player i is modeled a tree hierarchy: player i may know a set of distributions over types, a distribution over sets of distributions over types, or any other kinds of complex composition of uncertainty using “sets” and “distributions”.

Note that the Hierarchy Type Uncertainty model is a strict generalization of the all previous ones (i.e., exact, approximate, Bayesian and Knightian). We are unaware of any reasonable type uncertainty model (suitable for the setting of incomplete information) that cannot be defined within this model.

While a set-theoretic notion of Approximate Type Uncertainty is already very well motivated, we show that it is actually *universal* for modeling player uncertainties, at least for the very broad class of games with convex player types.

Theorem 7 (informal). *Whenever a player’s type is convex, the Hierarchy Type Uncertainty (and thus any exact, Bayesian, Knightian) model is equivalent to the Approximate Type Uncertainty model.*

This theorem also explains that why in this thesis we focus only on mechanism design with approximate types for auctions, which is just without loss of generality.

1.5 Related Work

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

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

In single-good auctions, Milgrom [23] studies the revenue difference between second-price and English auctions, when the players do not exactly know their valuations, but only that they are drawn from a *common* distribution.

Sandholm [30] presents an example of an auction (with an unconventionally-defined utility function) where a player’s valuation is drawn from the uniform distribution over $[0, 1]$, and argues that reporting the expected valuation (i.e., 0.5) is no longer dominant-strategy.

Porter et al. [28] consider a scheduling problem where tasks are to be assigned to players, and each player i privately knows that he would fail to perform task j with probability p_{ij} . This failure rate can be understood as a distribution of player’s private type. Their paper studies efficient dominant-strategy mechanisms in this setting.

Feige and Tennenholtz [14] consider the problem of scheduling n players to use the same machine. Each player i has a task requiring time length l_i , but he does not know l_i : he only knows that l_i is drawn from a distribution L_i . The authors study dominant-strategy mechanisms without monetary transfer, and prove that even if L_i ’s support has two elements, then no constant fraction of the maximum social welfare can be guaranteed. To overcome this difficulty, they introduce a different measure of social efficiency, which they call “fair share”, and provide mechanisms to guarantee an $\Omega(1)$ fair share.³

Set-theoretic models of “self uncertainty”. As already mentioned, in his work (further formalized by Bewley [5]) Knight [20] considers that a player has a set of distributions and knows that his true type is drawn from one of them. The Knightian model immediately implies that the preferences for a player are no longer *completely ordered*: some pairs of preferences may become incomparable. (For single-good auctions, as argued in Section 1.4, it is equivalent to each player i knowing a set of candidates K_i for his true valuation.)

Most of the works in the Knightian model address decision making. Some authors, such as Aumann [1], Dubra et al. [12], Ok [27] and Nascimento [25] work with incomplete orders. Others authors discuss various ways of bypassing the set-theoretic component of the Knightian model by computing a single number from the set of expected utilities: [11] picks the average or an arbitrary one, [31] picks the so-called Choquet expectation, and [17] picks the maximum.

General equilibrium models with unordered preferences have been considered by Mas-Colell [22], Gale and Mas-Colell [16], Shafer and Sonnenschein [32], and Fon and Otani [15]. More recently, Rigotti and Shannon [29] characterize the set of equilibria

³A “ ρ fair share” is a property such that each player i has at least ρ success rate if all other players share the same distribution as his l_i .

in a financial market problem.⁴

Lopomo and the previous two authors [21] also construct explicit mechanisms in a Knightian model, but for a single player. Specifically they consider the rent-extraction problem under two notions of implementation: 1) when reporting the truth is at least as good as any other strategy 2) when reporting the truth is not strictly eliminated in favor of another strategy. (Notice that, not envisaging other players, these are not notions of dominance, since the latter should take into account *all* strategy sub-profiles of other players.)

⁴ 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.

Chapter 2

Single-Good Auctions

We first choose to investigate the feasibility of mechanism design in the approximate-valuation model for a very simple and familiar application, maximizing social welfare in single-good auctions, for which the second-price mechanism gives us a simple, elegant, and perfect solution in the traditional, exact-valuation model.

2.1 The Model

As for any game, an auction of a single-good can be thought as consisting of two parts, a context and a mechanism. Later, to analyze players' behaviors in such an auction, we need solution concepts: implementation in dominant strategies and implementation in undominated strategies. As mentioned earlier, our specific goal in this chapter is to guarantee high social welfare as a performance measure. All these notions are going to be redefined in our approximate-valuation world.

2.1.1 The Auction Context

For $\delta \in [0, 1]$, a δ -approximate (auction) context C is a tuple $(n, B, \delta, \theta, K)$. In such a context there are n players, with all possible valuations being integers between 0 and a valuation bound B . Each player i does not know his own true valuation θ_i , but a set $K_i \subseteq \{0, \dots, B\}$ such that

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

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

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

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

There are two more components that are presumed in a single-good auction when a context \mathbf{C} is given:

- $\Omega = ([n] \cup \{\perp\}) \times \mathbb{R}^{[n]}$, the set of *outcomes*. If $(a, P) \in \Omega$ is an outcome, then we refer to a as its *allocation* and to P as its profile of *prices*. (If $a \in [n]$ then player a wins the good; if $a = \perp$ then the good remains unallocated.)
- Player i 's *utility* on outcome (a, P) is $U_i(\theta_i, (a, P))$, which is $\theta_i - P_i$ if $a = i$, and $-P_i$ otherwise.

Question and answer.

- **Why δ ?** In our purely set-theoretic framework, the approximation accuracy δ can be interpreted as quantifying the “quality of the players’ knowledge about themselves.” We thus find it natural to measure the performance of a mechanism as a function of δ . Without identifying any structure in the players’ possible multiple valuations, one may at most design elementary mechanisms, rather than “good” single-good auctions. In sum, the accuracy parameter δ is our Trojan horse for bringing meaningful mechanism design into our model.
- **Why $\theta_i \in K_i$?** A player’s approximate valuation contains the player’s true valuation because we consider each player to be “the ultimate authority about himself.”

- **Why is K_i not just an interval of integers?** An approximate valuation K_i may indeed be just that. But defining it to be a more general subset of $\{0, \dots, B\}$ may be necessary in some contexts. For instance, consider a player i who is about to participate to a yard-sale auction of a large sofa. He may know precisely the amount v he would pay for the sofa, but also that, if he wins it, he would have to carry it on top of his car, which is illegal and punishable with a fine f . In a Bayesian setting, he should compute his expected value of the sofa from the probability of being caught by the police on his way home (presumably based on the specific time of day, the immediate weather forecast, the immediate traffic forecast, the likelihood that other crimes might compete for the police's attention, and so on). But in our set-theoretic model, his approximate valuation K_i blissfully consists of just two separate values: namely, $K_i = \{v, v - f\}$.¹
- **Multiplicative or additive accuracy?** A greater level of generality is achieved by considering two distinct accuracy parameters: a multiplicative one, δ^* , and an additive one, δ^+ , leading to the following modified constraint:

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

All of our theorems hold for such more general approximate valuations. For simplicity, however, we consider only one kind of accuracy parameters, and we find the multiplicative one more meaningful.

- **Multiple possible δ 's?** Yes: indeed, $\delta > \delta'$ implies that every δ' -approximate context also is a δ -approximate one, and that $\mathcal{C}_{n,B}^\delta \supseteq \mathcal{C}_{n,B}^{\delta'}$.
- **Do the players know δ ?** A player i in a δ -approximate context may know nothing about a “global δ .” Of course, knowing K_i , he can certainly compute his smallest “local” δ_i : namely, $\frac{\max K_i - \min K_i}{\max K_i + \min K_i}$. But he may not have enough information about his opponents to realize that he is in a δ -approximate context for $\delta < 1$.
- **Does the designer know δ ?** When disproving the existence of mechanisms with a given efficiency guarantee for $\mathcal{C}_{n,B}^\delta$, we gladly assume that the mecha-

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

nism designer does know δ precisely, since this makes our impossibility results stronger. When proving the existence of such mechanisms, we shall specify whether or not the designer knows a “sufficient” δ .

- **Can real δ 's be really large?** Absolutely. Valuations may indeed be “very approximate.” Consider a firm participating to an auction for the exclusive rights to manufacture solar panels in the US for a period of 25 years. Even if the demand were precisely known in advance, and the only uncertainty were to come from the firm’s ability to lower its costs of production via some breakthrough research, an approximation accuracy of the firm’s own valuation for the license could easily exceed 0.5.

2.1.2 The Mechanism

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

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

NOTATION.

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

2.1.3 Solutions Concepts

In a non-Bayesian setting of incomplete information, two notions of implementations are natural to explore: implementation in dominant strategies and in undominated strategies [19]. However, the classical definitions of *dominance* need to be extended in order to apply to our approximate-valuation model. (Essentially, we adapt the generalized notions of dominance for Knightian uncertainty [5, 20] to our purely set-theoretic setting.)

Informally, a strategy s_i of a player i *very-weakly dominates* another strategy t_i of i , relative to a mechanism M and an approximate valuation K_i of i , if, for all candidate valuations in K_i and all possible strategy subprofiles of his opponents, s_i gives i a utility greater than or equal to that given him by t_i . If it is further the case that s_i gives i strictly greater utility than t_i for at least some valuations in K_i and strategy subprofiles of i 's opponents, then s_i *weakly dominates* t_i . A strategy of i is *undominated* if it is not weakly dominated. The set of such undominated strategies is denoted by $\text{UDed}_i(K_i)$ when M is clear from context, and let $\text{UDed}(K) \stackrel{\text{def}}{=} \text{UDed}_1(K_1) \times \cdots \times \text{UDed}_n(K_n)$. Notice that the following fact is obvious:

Fact 2.1. $\text{UDed}_i(K_i) \neq \emptyset$ for all K_i .

We defer the reasons of such choice and a more detailed definition to Chapter 7, when the approximate types for general games are introduced.

2.1.4 The Performance Measure

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

1. A context in $\mathcal{C}_{n,B}^\delta$ materializes: that is, there is one good for sale and n players show up, each player i with a δ -approximate valuation set K_i .

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

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

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

2.2 Our Results

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

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

2.2.1 The Inadequacy of Dominant-Strategy Mechanisms

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

This “reasoning” presupposes that, as in the exact-valuation world, an auction mechanism can safely restrict a player’s strategies to (reporting) single valuations. In our setting, however, it is not only reasonable but even natural for a mechanism to allow a player to report a *set of valuations* (e.g., his own K_i). Indeed, it is easy to realize that the revelation principle [24] continues to guarantee that every very-weakly-dominant-strategy mechanism for $\mathcal{C}_{n,B}^\delta$ has an equivalent *very-weakly-dominant-strategy-truthful* mechanism. In our approximate world, a mechanism is of the latter kind if, for each player i : (1) S_i , the strategies of i , consists of reporting an arbitrary δ -approximate set V_i of valuations, and (2) reporting his true approximate valuation K_i is a very-weakly-dominant strategy.

With such richer strategy sets, in principle there might be a dominant-strategy mechanism guaranteeing maximum social welfare. More realistically, in light of the approximate accuracy of the players’ self knowledge, one should expect some degradation of performance to be unavoidable. For instance, one might conjecture that, in a δ -approximate context, a dominant-strategy mechanism might be able to guarantee some δ -dependent fraction —such as $(1 - \delta)$, $(1 - 3\delta)$, or $(1 - \delta)^2$ — of the maximum social welfare. We prove, however, that also this is too optimistic.

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

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

(The proof of Theorem 1 can be found in Section 4.1.)

As a *relative* measure of the quality of the players’ self knowledge, δ should be

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

*dominant strategies are **intrinsically** linked to each player having exact²
knowledge of his own valuations.*

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

2.2.2 The Power of Deterministic Undominated-Strategy Mechanisms

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

²Here “exact” also includes exact Bayesian knowledge, which we will discuss in Chapter 8.

Theorem 2. *The following two statements hold:*

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

$$\text{SW}(\theta, F_{2P}(v)) \geq \left(\frac{1-\delta}{1+\delta} \right)^2 \text{MSW}(\theta) - 2 \frac{1-\delta}{1+\delta} .^3$$

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

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

(Theorem 2a and Theorem 2b are respectively proven in Section 4.4 and Section 4.2.⁴)

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

Easier and harder. Theorem 2a is not hard to prove. At a very high level, *“It is obvious that each player i should only consider bidding a value v_i inside his own approximate valuation K_i . It is further obvious that the worst possible gap between the maximum and the actual social welfare is achieved in the following case. Let w be the winner in the second-price mechanism, and let $h, h \neq w$, be the player with the largest candidate valuation. Player w bids $v_w = \max K_w$, and player h bids $v_h = \min K_h$ (and v_w only slightly exceeds v_h). In this case it is obvious that the second-price mechanism guarantees at most a fraction $\approx \left(\frac{1-\delta}{1+\delta} \right)^2$ of the maximum social welfare. \square ”*

³Breaking ties at random, the performance guarantee is only marginally better: namely, *exactly* $\left(\frac{1-\delta}{1+\delta} \right)^2 \text{MSW}(\theta)$.

⁴To denote a strategy profile, we use “ v ” in the statement of Theorem 2a to emphasize that each player in the M_{2P} indeed bids a valuation; and “ s ” in Theorem 2b to emphasize that a player’s strategies can be totally arbitrary.

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

Theorem 2b is harder to prove, as we should expect for all impossibility results. Working in undominated strategies, the revelation principle no longer applies. Thus, rather than analyzing a single mechanism (the “direct truthful” one), in principle we should consider *all* possible mechanisms, and establish that each one of them does no better than the second-price one. In particular, while the second-price mechanism has a clear and simple strategy space (namely, an integer in $\{0, 1, \dots, B\}$), we should consider mechanisms giving the players *absolutely arbitrary* strategies: even reporting arbitrary subsets of $\{0, 1, \dots, B\}$ would be a strong restriction! Establishing Theorem 2b thus requires new techniques, to be discussed in Section 3.1.

2.2.3 The Greater Power of Probabilistic Undominated-Strategy Mechanisms

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

Theorem 3. *The following two statements hold:*

- a. $\forall n, \forall \delta \in (0, 1)$, and $\forall B$, there exists a mechanism⁵ $M_{\text{opt}}^{(\delta)}$ such that for every δ -approximate-valuation profile K , every true-valuation profile $\theta \in K$, and every strategy profile $s \in \text{UDed}(K)$:

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

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

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

⁵See Chapter B for a concise definition of $M_{\text{opt}}^{(\delta)}$.

(The proofs of Theorem 3a and Theorem 3b can be found in Section 4.5 and Section 4.3 respectively. Note that the mechanism $M_{\text{opt}}^{(\delta)}$ of Theorem 3a is constructed given knowledge of δ .⁶)

Theoretical significance. Theorem 3 highlights a novelty of the approximate-valuation world: namely, probabilism enhances the power of implementation in undominated strategies even for guaranteeing social welfare. By contrast, probabilism offers no such advantage in the exact-valuation world, since the deterministic second-price mechanism already guarantees maximum social welfare. We conjecture that, in the approximate-valuation world, probabilistic mechanisms will enjoy a provably better performance in other applications as well.

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

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

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

⁶With extra pains, however, one can actually get a reasonable performance if the mechanism only knows a good upper bound on δ .

Chapter 3

Single-Good Auctions:

Techniques of Independent Interest

New ventures require new tools. We thus wish to highlight two techniques that we believe will prove useful to the design and analysis of mechanisms in the approximate valuation setting. We shall prove them in the setting of a single-good auction, but notice that it is a simple exercise to have the *Undominated Intersection Lemma* generalized to any game setting, and the *Distinguishable Monotonicity Lemma* generalized to any single-parameter setting. (See Chapter 7 for notions of our approximate types in general games.)

3.1 The Undominated Intersection Lemma

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

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

$$\text{UDed}_i(K_i) = K_i \quad ? \tag{3.1}$$

If so, this would imply the following:

$$K_i \cap \tilde{K}_i \neq \emptyset \Rightarrow \text{UDed}_i(K_i) \cap \text{UDed}_i(\tilde{K}_i) \neq \emptyset . \quad (3.2)$$

However, relation (3.1) is in general false. Even if the strategies given to each player are valuations between 0 and B , a mechanism does not need to interpret a bid v_i reported by i as i 's true valuation θ_i . For instance, the mechanism could first replace each v_i by $\pi(v_i)$ where π is some fixed permutation over $\{0, 1, \dots, B\}$ and then run the second-price mechanism as if each player i had bid $\pi(v_i)$. In this case, after $\text{UDed}(K_i)$ has been correctly computed, it will look very different from K_i .

Relation (3.2) might hold even if relation (3.1) does not. But it is unclear that it does: the set of strategies S_i 's provided by a mechanism can be absolutely arbitrary, rather than $\{0, 1, \dots, B\}$. Therefore, as soon as K_i and \tilde{K}_i are even slightly different, their corresponding $\text{UDed}_i(K_i)$ and $\text{UDed}_i(\tilde{K}_i)$ may in principle be totally unrelated.

We prove however that *a sufficiently simple variant of relation (3.2) holds for all mechanisms*, not just the ones with $S_i = \{0, 1, \dots, B\}$. Informally,

For any mechanism, no matter whether it is probabilistic or not, if K_i and \tilde{K}_i have at least two values in common, then there exist two (possibly mixed) "almost payoff-equivalent" strategies σ_i and $\tilde{\sigma}_i$ respectively having $\text{UDed}_i(K_i)$ and $\text{UDed}_i(\tilde{K}_i)$ as their support.^a

^aActually relation (3.2) holds when the total number of coins usable by the players is bounded.

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

Let us emphasize that this lemma actually applies to *all* undominated-strategy mechanisms in the approximate-valuation world, not just to auction mechanisms, let alone single-good ones. The *Undominated Intersection Lemma* is the key to our impossibility results for undominated strategies.

3.1.1 Details

Lemma 3.1 (Undominated Intersection Lemma). *Let $M = (S, F)$ be a mechanism, i a player, and K_i and \tilde{K}_i two approximate valuations of i intersecting in at least two integers. Then, for every $\varepsilon > 0$, there exist mixed strategies $\sigma_i \in \Delta(\text{UDed}_i(K_i))$ and $\tilde{\sigma}_i \in \Delta(\text{UDed}_i(\tilde{K}_i))$ such that*

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

Proof. Let x_i and y_i be two distinct integers in $K_i \cap \tilde{K}_i$, and, without loss of generality, let $x_i > y_i$. Recall that, by Fact 2.1, $\text{UDed}_i(K_i)$ and $\text{UDed}_i(\tilde{K}_i)$ are both nonempty.

If there exists a common (pure) strategy $s_i \in \text{UDed}_i(K_i) \cap \text{UDed}_i(\tilde{K}_i)$, then setting $\sigma_i = \tilde{\sigma}_i = s_i$ completes the proof. Therefore, let us assume that $\text{UDed}_i(K_i)$ and $\text{UDed}_i(\tilde{K}_i)$ are disjoint, and let s_i be a strategy in $\text{UDed}_i(K_i)$ but not in $\text{UDed}_i(\tilde{K}_i)$. The finiteness of the strategy set S_i implies the existence of a strategy $\tilde{\sigma}_i \in \Delta(\text{UDed}_i(\tilde{K}_i))$ such that $\tilde{\sigma}_i \succ_{i, \tilde{K}_i}^w s_i$. We now argue that

$$\exists \tau_i \in \Delta(\text{UDed}_i(K_i)) \text{ such that } \tau_i \succ_{i, K_i}^w \tilde{\sigma}_i. \quad (3.3)$$

Letting $\tilde{\sigma}_i = \sum_{j \in X} \alpha^{(j)} s_i^{(j)}$ —where X is a subset of S_i —and invoking again the disjointness of the two undominated strategy sets, we deduce that for each $j \in X$ there exists a strategy $\tau_i^{(j)} \in \Delta(\text{UDed}_i(K_i))$ such that $\tau_i^{(j)} \succ_{i, K_i}^w s_i^{(j)}$. Thus, $\tau_i \stackrel{\text{def}}{=} \sum_{j \in X} \alpha^{(j)} \tau_i^{(j)}$ satisfies Eq. (3.3).

For the same reason, we can also find some $\tilde{\tau}_i \in \Delta(\text{UDed}_i(\tilde{K}_i))$ such that $\tilde{\tau}_i \succ_{i, \tilde{K}_i}^w \tau_i$. Continuing in this fashion, “jumping” back and forth between $\Delta(\text{UDed}_i(K_i))$ and $\Delta(\text{UDed}_i(\tilde{K}_i))$, we obtain an infinite chain of (not necessarily distinct) strategies, $\{\sigma_i^{(k)}, \tilde{\sigma}_i^{(k)}\}_{k \in \mathbb{N}}$, where:

$$\sigma_i^{(1)} \succ_{i, K_i}^w \tilde{\sigma}_i^{(1)} \succ_{i, \tilde{K}_i}^w \sigma_i^{(2)} \succ_{i, K_i}^w \tilde{\sigma}_i^{(2)} \succ_{i, \tilde{K}_i}^w \dots$$

¹Note that, while we have only defined what it means for a *pure* strategy to be dominated by a possibly mixed one, the definition trivially extends to the case of dominated strategies that are *mixed*, as is the case in “ $\tau_i \succ_{i, K_i}^w \tilde{\sigma}_i$ ” in Eq. (3.3).

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

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

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

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

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

$$\begin{aligned} F_i^A(\sigma_i^{(1)} \sqcup s_{-i})x_i - F_i^P(\sigma_i^{(1)} \sqcup s_{-i}) &\leq F_i^A(\tilde{\sigma}_i^{(2)} \sqcup s_{-i})x_i - F_i^P(\tilde{\sigma}_i^{(2)} \sqcup s_{-i}) \\ &\leq F_i^A(\sigma_i^{(3)} \sqcup s_{-i})x_i - F_i^P(\sigma_i^{(3)} \sqcup s_{-i}) \leq \dots \end{aligned}$$

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

$$F_i^A(\sigma_i^{(k)} \sqcup s_{-i})x_i - F_i^P(\sigma_i^{(k)} \sqcup s_{-i}) \leq F_i^A(\tilde{\sigma}_i^{(k)} \sqcup s_{-i})x_i - F_i^P(\tilde{\sigma}_i^{(k)} \sqcup s_{-i}) \quad (3.4)$$

$$F_i^A(\tilde{\sigma}_i^{(k)} \sqcup s_{-i})x_i - F_i^P(\tilde{\sigma}_i^{(k)} \sqcup s_{-i}) \leq F_i^A(\sigma_i^{(k+1)} \sqcup s_{-i})x_i - F_i^P(\sigma_i^{(k+1)} \sqcup s_{-i}) \quad (3.5)$$

$$F_i^A(\sigma_i^{(k+1)} \sqcup s_{-i})x_i - F_i^P(\sigma_i^{(k+1)} \sqcup s_{-i}) \leq F_i^A(\sigma_i^{(k)} \sqcup s_{-i})x_i - F_i^P(\sigma_i^{(k)} \sqcup s_{-i}) + \frac{\varepsilon}{3B} . \quad (3.6)$$

Similarly, because $y_i \in K_i \cap \tilde{K}_i$, setting $\theta_i = \tilde{\theta}_i = y_i$, we have that for every

$s_{-i} \in S_{-i}$ there must exist some $H_\varepsilon^{(s_{-i}, y_i)} \in \mathbb{N}$ such that $\forall k > H_\varepsilon^{(s_{-i}, y_i)}$:

$$F_i^A(\sigma_i^{(k)} \sqcup s_{-i})y_i - F_i^P(\sigma_i^{(k)} \sqcup s_{-i}) \leq F_i^A(\tilde{\sigma}_i^{(k)} \sqcup s_{-i})y_i - F_i^P(\tilde{\sigma}_i^{(k)} \sqcup s_{-i}) \quad (3.7)$$

$$F_i^A(\tilde{\sigma}_i^{(k)} \sqcup s_{-i})y_i - F_i^P(\tilde{\sigma}_i^{(k)} \sqcup s_{-i}) \leq F_i^A(\sigma_i^{(k+1)} \sqcup s_{-i})y_i - F_i^P(\sigma_i^{(k+1)} \sqcup s_{-i}) \quad (3.8)$$

$$F_i^A(\sigma_i^{(k+1)} \sqcup s_{-i})y_i - F_i^P(\sigma_i^{(k+1)} \sqcup s_{-i}) \leq F_i^A(\sigma_i^{(k)} \sqcup s_{-i})y_i - F_i^P(\sigma_i^{(k)} \sqcup s_{-i}) + \frac{\varepsilon}{3B} . \quad (3.9)$$

Notice now that, because the set of strategies S_{-i} is finite,

$$H_\varepsilon = \max_{s_{-i} \in S_{-i}} \{H_\varepsilon^{(s_{-i}, x_i)}, H_\varepsilon^{(s_{-i}, y_i)}\}$$

is a well defined integer. Next, we pick arbitrarily $\bar{k} > H_\varepsilon$, and prove that $\sigma_i^{(\bar{k}+1)}$ and $\tilde{\sigma}_i^{(\bar{k})}$ are the two strategies that we are looking for.

To this end, pick arbitrarily $s_{-i} \in S_{-i}$ and sum up Eq. (3.4), Eq. (3.6) and Eq. (3.8). The (expected) prices and the $F_i^A(\sigma_i^{(\bar{k})}, s_{-i})x_i$ term will cancel out so as to yield

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

Then, sum up Eq. (3.5), Eq. (3.7) and Eq. (3.9). The (expected) prices and the $F_i^A(\sigma_i^{(\bar{k})}, s_{-i})y_i$ term will cancel out yielding

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

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

$$|F_i^A(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i}) - F_i^A(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i})| \leq \frac{\varepsilon}{3B} \leq \varepsilon . \quad (3.10)$$

That is, the first inequality of Lemma 3.1 has been proved. Let us now consider the price terms.

Fixing arbitrarily $s_{-i} \in S_{-i}$ and combining Eq. (3.5) and Eq. (3.10), we get:

$$\begin{aligned}
F_i^A(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i})x_i - F_i^P(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i}) &\leq F_i^A(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i})x_i - F_i^P(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i}) \\
&\leq \left(F_i^A(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i}) + \frac{\varepsilon}{3B} \right) x_i - F_i^P(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i}) \\
&\Rightarrow -F_i^P(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i}) \leq \frac{\varepsilon}{3} - F_i^P(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i}) . \quad (3.11)
\end{aligned}$$

Summing up Eq. (3.4) and Eq. (3.6) and then substituting Eq. (3.10), we get:

$$\begin{aligned}
F_i^A(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i})x_i - F_i^P(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i}) &\leq F_i^A(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i})x_i - F_i^P(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i}) + \frac{\varepsilon}{3B} \\
&\leq \left(F_i^A(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i}) + \frac{\varepsilon}{3B} \right) x_i - F_i^P(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i}) + \frac{\varepsilon}{3B} \\
&\Rightarrow -F_i^P(\sigma_i^{(\bar{k}+1)} \sqcup s_{-i}) \leq \frac{2\varepsilon}{3} - F_i^P(\tilde{\sigma}_i^{(\bar{k})} \sqcup s_{-i}) . \quad (3.12)
\end{aligned}$$

Finally, combining inequalities Eq. (3.11) and Eq. (3.12) we immediately get that for all $s_{-i} \in S_{-i}$:

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

That is, the second desired inequality also holds, completing the proof of Lemma 3.1. \square

3.2 The Distinguishable Monotonicity Lemma

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

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

Specifically, we consider mechanisms whose strategies consist of possible valuations, namely the set $\{0, \dots, B\}$, and whose allocation functions are restrictions (to

$\{0, 1, \dots, B\}^N$) of integrable functions (over $[0, B]^N$) satisfying a suitable monotonicity property. A simple lemma, the *Distinguishable Monotonicity Lemma*, will then guarantee that

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

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

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

3.2.1 Details

Before we describe our lemma, let us recall a traditional way to define auction mechanisms from suitable allocation functions.

Definition 3.2. *If $f : [0, B]^{[n]} \rightarrow [0, 1]^{[n]}$ is an integrable² allocation function, then we denote by M_f the mechanism (S, F) where $S = \{0, 1, \dots, B\}^{[n]}$ and F is so defined: on input bid profile $v \in S$,*

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

²Specifically, we require that, for each v_{-i} , the function $f_i(z \sqcup v_{-i})$ is integrable with respect to z on $[0, B]$.

Remark 3.3.

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

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

Definition 3.4. Let $f : [0, B]^{[n]} \rightarrow [0, 1]^{[n]}$ be a allocation function. For $d \in \{1, 2\}$, we say that f is **d -distinguishably monotonic** (d -DM, for short) if f is integrable, monotonic, and satisfying the following “distinguishability” condition:

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

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

Distinguishable monotonicity is certainly an additional requirement to monotonicity, but actually quite mild. Indeed, the second-price mechanism is 2-DM and, if ties are broken at random, even 1-DM (see Example 3.5). Yet, in our approximate-valuation world, this mild additional requirement is quite useful for “controlling” the *undominated strategies* of a mechanism, and thus for engineering implementations of desirable social choice functions in undominated strategies.

Example 3.5 (Second-Price Mechanism). Recall that the *second-price mechanism* is a direct mechanism that assigns the good to the highest bidder at a price equal to the second-highest bid: it is a pair $M_{2P} = (S_{2P}, F_{2P})$ with $S_{2P} = \{0, 1, \dots, B\}^{[n]}$ and

$F_{2P}(v) \stackrel{\text{def}}{=}$

1. Assign the good to the highest bidder: $i^* \stackrel{\text{def}}{=} \arg \max_{i \in [n]} v_i$.
2. Charge the highest bidder the second price: $P_{i^*} \stackrel{\text{def}}{=} \max_{i \in [n], i \neq i^*} v_i$.

(And everyone else pays nothing.)

In the language of allocation functions, M_{2P} can be represented as a mechanism M_f , where f is 2-DM; also, if we require that the ties (for the highest bidder) are broken at random (by giving a positive, but not necessarily equal, probability to every highest bidder), then f is 1-DM. So let us prove these two facts:

Proof. At a high level, the allocation function f for M_{2P} is almost unique, except for those input bids that contain ties. Now take an arbitrary second-price mechanism M_{2P} with a specific tie breaking rule. For each player $i \in [n]$, we define f_i as follows: for every bid sub-profile $z_{-i} \in [0, B]^{[n]-\{i\}}$, letting $x^* \stackrel{\text{def}}{=} \max_{j \neq i} z_j$,

- for every $x < x^*$, define $f_i(x \sqcup z_{-i}) \stackrel{\text{def}}{=} 0$,
- for every $x > x^*$, define $f_i(x \sqcup z_{-i}) \stackrel{\text{def}}{=} 1$,
- for $x = x^*$ then there is a tie, in which case
 - if z_{-i} is a valid integer bid in $\{0, 1, \dots, B\}^{[n]-\{i\}}$, then define $f_i(x^* \sqcup z_{-i})$ to be the winning probability according to how M_{2P} breaks the tie,³ and
 - if z_{-i} is not a valid integer bid in $\{0, 1, \dots, B\}^{[n]-\{i\}}$, then define $f_i(x^* \sqcup z_{-i})$ arbitrarily (say 0 for example).⁴

One can verify that the mechanism M_f according to the definition above is exactly the given M_{2P} ; this is because the two coincide on the allocation probabilities for all integer points $v \in \{0, 1, \dots, B\}^{[n]}$, and the price (recall the integral in Definition 3.2) is exactly the winning probability multiplied by the second highest bid.

It is clear that f is monotonic. Moreover, for any $v_i, v'_i \in \{0, 1, \dots, B\}$ such that $v_i < v'_i - 1$, let everyone else bid some integer x in the open interval (v_i, v'_i) by setting $v_{-i} = \{x, x, \dots, x\}$. By construction, $f_i(v_i \sqcup v_{-i}) = 0$ and every (not necessarily

³If in M_{2P} player i receives the good with probability 1, then we set $f_i(x^* \sqcup z_{-i}) \stackrel{\text{def}}{=} 1$; if player i receives the good with probability 0.2, then $f_i(x^* \sqcup z_{-i}) \stackrel{\text{def}}{=} 0.2$, and so on.

⁴Indeed, M_f will never be invoked on an input with more than one non-integer points. It invokes integer points for calculating allocation probabilities, and one non-integer points for calculating the price.

integer) $z \in (x, v'_i]$ satisfies $f(z \sqcup v_{-i}) = 1$; but this establishes that f is 2-DM:

$$\int_{v_i}^{v'_i} (f_i(z \sqcup v_{-i}) - f_i(v_i \sqcup v_{-i})) dz \geq (v'_i - x)(1 - 0) \geq 1 > 0 .$$

If instead we tweak M_{2P} to break ties at random by giving a positive (but not necessarily equal) probability to every highest bidder, then, for any $v_i, v'_i \in \{0, 1, \dots, B\}$ such that $v_i < v'_i$, let everyone else bid some $x = v_i$ by setting $v_{-i} = \{x, x, \dots, x\}$. We have that for every $z \in (x, v'_i]$, $f(z \sqcup v_{-i}) = 1$, but $f(v_i \sqcup v_{-i}) < 1$ (since every highest bidder is awarded the good with positive probability); but this establishes that f is 1-DM:

$$\int_{v_i}^{v'_i} (f_i(z \sqcup v_{-i}) - f_i(v_i \sqcup v_{-i})) dz \geq (v'_i - x) \cdot (1 - f(v_i \sqcup v_{-i})) > 0 ,$$

as desired. □

Lemma 3.6 (Distinguishable Monotonicity Lemma). *If f is a d -DM allocation function, then M_f is such that, for any player i and δ -approximate-valuation profile K ,*

$$\text{UDed}_i(K_i) \subseteq \{\min K_i, \dots, \max K_i\} \text{ if } d = 1, \text{ and}$$

$$\text{UDed}_i(K_i) \subseteq \{\min K_i - 1, \dots, \max K_i + 1\} \text{ if } d = 2.$$

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

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

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

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

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

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

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

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

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

Chapter 4

Single-Good Auctions:

Missing Proofs

4.1 Proof of Theorem 1

Recall that, with the help from the Revelation Principle (which we will prove independently in Section 7.5), it suffices to prove this negative result for “direct” mechanisms where players report arbitrary δ -approximate sets of valuations, and for each player reporting his true approximate valuation K_i is a very-weakly-dominant strategy.

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

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

Proof. Fix arbitrarily n , δ , and B such that $B > \frac{3-\delta}{2\delta}$. We start by proving a separate claim: essentially, if a player reports a δ -interval whose center is sufficiently high, then his winning probability and expected price remain constant.

Claim 4.1. *For all players i , integers $x \in (\frac{3-\delta}{2\delta}, B]$, and δ -approximate-valuation sub-profiles \tilde{K}_{-i} ,*

$$F_i^A(\delta[x] \sqcup \tilde{K}_{-i}) = F_i^A(\delta[x+1] \sqcup \tilde{K}_{-i}) \tag{4.1}$$

and

$$F_i^P(\delta[x] \sqcup \tilde{K}_{-i}) = F_i^P(\delta[x+1] \sqcup \tilde{K}_{-i}) . \quad (4.2)$$

Proof of Claim 4.1. Because the approximate valuation K_i of player i can be $\delta[x]$, and because when this is the case reporting the truth $\delta[x]$ very-weakly dominates $\delta[x+1]$, the following inequality must hold: $\forall \theta_i \in \delta[x]$,

$$F_i^A(\delta[x] \sqcup \tilde{K}_{-i}) \cdot \theta_i - F_i^P(\delta[x] \sqcup \tilde{K}_{-i}) \geq F_i^A(\delta[x+1] \sqcup \tilde{K}_{-i}) \cdot \theta_i - F_i^P(\delta[x+1] \sqcup \tilde{K}_{-i}) \quad (4.3)$$

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

$$F_i^A(\delta[x+1] \sqcup \tilde{K}_{-i}) \cdot \theta'_i - F_i^P(\delta[x+1] \sqcup \tilde{K}_{-i}) \geq F_i^A(\delta[x] \sqcup \tilde{K}_{-i}) \cdot \theta'_i - F_i^P(\delta[x] \sqcup \tilde{K}_{-i}) . \quad (4.4)$$

Thus, setting $\theta_i = x$ in Eq. (4.3) and $\theta'_i = x+1$ in Eq. (4.4), and summing up the resulting inequalities, the F_i^P price terms and a few other terms cancel out yielding the following inequality:

$$F_i^A(\delta[x+1] \sqcup \tilde{K}_{-i}) \geq F_i^A(\delta[x] \sqcup \tilde{K}_{-i}) . \quad (4.5)$$

Also, setting $\theta_i = \lfloor x(1+\delta) \rfloor$ in Eq. (4.3) and $\theta'_i = \lceil (x+1)(1-\delta) \rceil$ in Eq. (4.4),¹ and summing up the resulting inequalities we obtain the following one:

$$\left(F_i^A(\delta[x] \sqcup \tilde{K}_{-i}) - F_i^A(\delta[x+1] \sqcup \tilde{K}_{-i}) \right) \cdot \left(\lfloor x(1+\delta) \rfloor - \lceil (x+1)(1-\delta) \rceil \right) \geq 0 . \quad (4.6)$$

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

Thus from Eq. (4.6) we deduce

$$F_i^A(\delta[x] \sqcup \tilde{K}_{-i}) \geq F_i^A(\delta[x+1] \sqcup \tilde{K}_{-i}) \quad (4.7)$$

Together, Eq. (4.5) and Eq. (4.7) imply the desired Eq. (4.1). Finally, combining Eq. (4.1) with Eq. (4.3) and Eq. (4.4) we obtain the desired Eq. (4.2). \square

Let us now finish the proof of Theorem 1.

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

Choose the profile of approximate valuations $\widehat{K} \stackrel{\text{def}}{=} (\delta[c], \delta[c], \dots, \delta[c])$, where $c \stackrel{\text{def}}{=} \lfloor \frac{3-\delta}{2\delta} \rfloor + 1$. By averaging, because the summation of $F_i^A(\widehat{K})$ over $i \in [n]$ cannot be greater than 1, there must exist a player j such that $F_j^A(\widehat{K}) \leq 1/n$. Without loss of generality, let such player be player 1. Then, invoking Claim 4.1 multiple times we have

$$\begin{aligned} F_1^A(\delta[B], \delta[c], \dots, \delta[c]) &= F_1^A(\delta[B-1], \delta[c], \dots, \delta[c]) = \dots \\ &= F_1^A(\delta[c], \delta[c], \dots, \delta[c]) = F_1^A(\widehat{K}) \leq \frac{1}{n}. \end{aligned}$$

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

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

as desired. ■

4.2 Proof of Theorem 2b

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

$$\text{SW}(\theta, F(s)) \leq \left(\left(\frac{1-\delta}{1+\delta} \right)^2 + \frac{3}{B} \right) \text{MSW}(\theta). \quad (4.8)$$

Proof. Choose $x \stackrel{\text{def}}{=} \frac{B}{1+\delta}$ and $y \stackrel{\text{def}}{=} \frac{(1-\delta)x+2}{1+\delta}$, we have $x \geq y$ due to our choice of $B \geq \frac{1+\delta}{\delta}$. Recalling that $\delta[x] \stackrel{\text{def}}{=} [(1-\delta)x, (1+\delta)x] \cap \{0, 1, \dots, B\}$, one can verify that $\delta[x]$ and $\delta[y]$ both contain the two integers $\lceil (1-\delta)x \rceil$ and $\lceil (1-\delta)x \rceil + 1$,² satisfying the requirement of the Undominated Intersection Lemma 3.1.

Choose ε such that $\frac{1}{n} + \varepsilon < 1$. Then the Undominated Intersection Lemma 3.1

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

guarantees that

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

$$|F_i^A(\sigma_i \sqcup s_{-i}) - F_i^A(\sigma'_i \sqcup s_{-i})| < \varepsilon . \quad (4.9)$$

Now consider the allocation distribution $F^A(\sigma'_1, \dots, \sigma'_n)$, where the randomness comes from the mixed strategy profile since M is a deterministic mechanism. Since the good will be assigned with a total probability mass of 1, by averaging, there exists a player j such that $F_j^A(\sigma'_1, \dots, \sigma'_n) \leq \frac{1}{n}$: that is, player j wins the good with probability at most $\frac{1}{n}$. Without loss of generality, let $j = 1$. In particular, there exist $s'_{-1} \in \text{UDed}_2(\delta[y]) \times \dots \times \text{UDed}_n(\delta[y])$, such that $F_1^A(\sigma'_1, s'_{-1}) \leq \frac{1}{n}$. This together with Eq. (4.9) implies that $F_1^A(\sigma_1 \sqcup s'_{-1}) \leq \frac{1}{n} + \varepsilon < 1$. In turn, this implies that there exists a *pure* strategy $s_1 \in \text{UDed}_1(\delta[x])$ such that, setting $s \stackrel{\text{def}}{=} (s_1 \sqcup s'_{-1})$, $F_1^A(s) = 0$.

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

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

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

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

Thus the theorem holds. ■

4.3 Proof of Theorem 3b

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

strategy profile $s \in \text{UDed}(K)$, and a true-valuation profile $\theta \in K$ such that

$$\mathbb{E}[\text{SW}(\theta, F(s))] \leq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} + \frac{3}{B} \right) \text{MSW}(\theta) . \quad (4.10)$$

Proof. (The first part of the proof closely tracks that of Theorem 2b in Section 4.2.³)

Choose again $x \stackrel{\text{def}}{=} \frac{B}{1+\delta}$ and $y \stackrel{\text{def}}{=} \frac{(1-\delta)x+2}{1+\delta}$, we have $x \geq y$ due to our choice of $B \geq \frac{1+\delta}{\delta}$, and $\delta[x]$ and $\delta[y]$ both contain the two integers $\lceil (1-\delta)x \rceil$ and $\lceil (1-\delta)x \rceil + 1$, satisfying the requirement of the Undominated Intersection Lemma 3.1.

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

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

Then the Undominated Intersection Lemma 3.1 guarantees that

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

$$|F_i^A(\sigma_i \sqcup s_{-i}) - F_i^A(\sigma'_i \sqcup s_{-i})| < \varepsilon . \quad (4.11)$$

Again consider the allocation distribution $F^A(\sigma'_1, \dots, \sigma'_n)$. By averaging, there exists some player j such that $F_j^A(\sigma'_1, \dots, \sigma'_n) \leq \frac{1}{n}$. Without loss of generality let $j = 1$. Thus, by our choice of ε and Eq. (4.11), we have that $F_1^A(\sigma_1 \sqcup \sigma'_{-1}) \leq \frac{1}{n} + \varepsilon$. This implies that there exists a *pure* strategy profile $s = (s_1 \sqcup s'_{-1})$ that is in the support of $(\sigma_1 \sqcup \sigma'_{-1})$ —and thus in $\text{UDed}_1(\delta[x]) \times \text{UDed}_2(\delta[y]) \times \dots \times \text{UDed}_2(\delta[y])$ —such that $F_1^A(s_1 \sqcup s'_{-1}) \leq \frac{1}{n} + \varepsilon$. Now define

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

Notice that $s \in \text{UDed}(K)$, $\theta \in K$, and $\text{MSW}(\theta) = (1+\delta)x = B$. We now show

³Very informally, the only differences are that the allocation distribution $F^A(\sigma'_1, \dots, \sigma'_n)$ now depends also on the “coin tosses of the mechanism”, and that one can no longer guarantee the existence of a pure strategy s such that $F_1^A(s) = 0$.

that s , K , and θ satisfy the desired Eq. (4.10):

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

■

4.4 Proof of Theorem 2a

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

$$\text{SW}(\theta, F_{2P}(v)) \geq \left(\frac{1 - \delta}{1 + \delta}\right)^2 \text{MSW}(\theta) - 2 \frac{1 - \delta}{1 + \delta} . \quad (4.12)$$

Proof. Let $M_{2P} = (S_{2P}, F_{2P})$ be (a version of) the second-price mechanism with a deterministic tie-breaking rule. Since K is a δ -approximate valuation, for each player i let x_i be such that $K_i \subseteq \delta[x_i] \cap \{0, \dots, B\}$. Then, in light of the Distinguishably Monotonicity Lemma 3.6 and the previous observation that F_{2P}^A is a 2-DM allocation function, we have that, for each player i :

$$\text{UDed}_i(x) \subseteq \{\lceil(1 - \delta)x_i\rceil - 1, \dots, \lfloor(1 + \delta)x_i\rfloor + 1\} . \quad (4.13)$$

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

$v_{j^*} = \max_j v_j$.) We now bound the difference between θ_{i^*} and θ_{j^*} when $i^* \neq j^*$.

From Eq. (4.13) we know that $\lceil (1 - \delta)x_{i^*} \rceil - 1 \leq v_{i^*}$ and $v_{j^*} \leq \lfloor (1 + \delta)x_{j^*} \rfloor + 1$. Because j^* is the winner, we also know that $v_{i^*} \leq v_{j^*}$. Combining these facts and removing “floors and ceilings” we have $(1 - \delta)x_{i^*} \leq (1 + \delta)x_{j^*} + 2$. Since we also know that $\theta_{j^*} \geq (1 - \delta)x_{j^*}$ and $(1 + \delta)x_{i^*} \geq \theta_{i^*}$, we obtain:

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

Thus, the claim of our theorem holds. \blacksquare

Consider the case where the second-price mechanism breaks ties at random (assigning a positive probability to each tie). Then, one can use a proof analogous to the one above, with the only difference being that F_{2P}^A is 1-DM (instead of only 2-DM), and invoking the stronger inclusion of the Distinguishably Monotonicity Lemma 3.6, to show the following, stronger lower bound:

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

4.5 Proof of Theorem 3a

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

4.5.1 A Very Restricted Search

In order to leverage our Distinguishably Monotonicity Lemma 3.6, it is natural for us to search for $M_{\text{opt}}^{(\delta)}$ among 1-DM mechanisms. Let us now distill an additional requirement for the underlying allocation function of such mechanisms that suffices for our goals. We shall do so in terms of the following positive quantity D_δ : for all

$\delta \in (0, 1)$,

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

Definition 4.2. We say that a allocation function f is δ -good if it is 1-DM and:

$$\forall i \in [n], \forall v \in \{0, 1, \dots, B\}^{[n]}, \quad \sum_{j=1}^n f_j(v)v_j + D_\delta \cdot f_i(v)v_i \geq \frac{1}{n} \cdot v_i(n + D_\delta) . \quad (4.14)$$

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

Lemma 4.3. If f is δ -good, then M_f satisfies that such that for every δ -approximate-valuation profile K , every strategy profile $s \in \text{UDed}(K)$ and every true-valuation profile $\theta \in K$:

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

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

$$\forall \theta \in K, \quad \forall v \in \text{UDed}(K), \quad \forall i \in [n], \quad \sum_{j=1}^n \theta_j f_j(v) \geq \left(\frac{(1 - \delta)^2 + \frac{4\delta}{n}}{(1 + \delta)^2} \right) \theta_i . \quad (4.15)$$

For every $i \in [n]$, let $x_i \in \mathbb{R}$ be such that $K_i \subseteq \delta[x_i]$, and let $\delta[x] = \delta[x_1] \times \dots \times \delta[x_n]$. Then, $\theta \in K$ and the Distinguishable Monotonicity Lemma respectively imply

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

Combining these two chains of inequalities yields

$$\frac{1 - \delta}{1 + \delta} v_i \leq \theta_i \leq \frac{1 + \delta}{1 - \delta} v_i . \quad (4.16)$$

Let us now argue that Eq. (4.15) holds by arbitrarily fixing v and i and showing that it is impossible to construct a “bad” θ so as to violate Eq. (4.15).

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

Considering these extreme choices, we conclude that no θ contradicts Eq. (4.15)

if:

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

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

Simplifying the above equations, Eq. (4.15) holds if both the following inequalities hold:

$$\sum_{j=1}^n v_j f_j(v) \geq \frac{n + D_\delta}{n} \cdot \frac{1}{D_\delta + 1} \cdot v_i, \quad (4.17)$$

$$\sum_{j=1}^n v_j f_j(v) + D_\delta \cdot v_i f_i(v) \geq \frac{n + D_\delta}{n} v_i. \quad (4.18)$$

Note that Eq. (4.18) holds because it is implied by the hypothesis that f is δ -good; note also that Eq. (4.17) holds because it is implied by Eq. (4.18). Indeed, since $\frac{1}{D_\delta + 1} = \left(\frac{1-\delta}{1+\delta}\right)^2 < 1$ for all $\delta \in (0, 1)$,

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

Thus both Eq. (4.15) and our lemma hold. \square

4.5.2 Our Allocation Function

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

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

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

Recall that, to be δ -good, a allocation function $f : [0, B]^{[n]} \rightarrow [0, 1]^{[n]}$ should satisfy Eq. (4.14), that is:

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

A reasonable guess to “solve for f ” is to restrict our attention to *symmetric* functions.

The most natural candidate is simply

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

One could verify that the function f , in addition to being symmetric, sums up to 1, is 1-DM, and satisfies the desired condition Eq. (4.14). (In fact, as we shall see, the above candidate f coincides with our proposed $f^{(\delta)}$ when no threshold is introduced.) We would be done, except for one crucial fact: f sometimes takes negative values!

We therefore need to “patch” the guessed function f by forcing non-negativity while maintaining the other required properties, and this is exactly where the idea of a threshold, winners, and losers comes in. Roughly, only players with sufficiently low reported valuations are at risk of a “negative probability” and, because are most likely to have low true valuations, we remove them from the auction altogether. To preserve the other properties, though, we need to re-weight the function, thereby obtaining Eq. (4.19). Thus, at high level, we simply keep removing players until all of the players are given non-negative probability (by virtue of being in the auction or having been thrown out). A similar idea previously appeared in [6].

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

varies too.)

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

Definition 4.4. For every $\delta \in (0, 1)$, define the function $f^{(\delta)}: [0, B]^{[n]} \rightarrow [0, 1]^{[n]}$ as follows: for every $i \in [n]$ and every $z = (z_1, \dots, z_n) \in [0, B]^{[n]}$

- if $z_1 \geq z_2 \geq \dots \geq z_n$, then

$$f_i^{(\delta)}(z) \stackrel{\text{def}}{=} \begin{cases} \frac{1}{n} \cdot \frac{n+D_\delta}{n^*+D_\delta} \cdot \frac{z_i(n^*+D_\delta) - \sum_{j=1}^{n^*} z_j}{z_i D_\delta}, & \text{if } i \leq n^*, \\ 0, & \text{if } i > n^*; \end{cases} \quad (4.19)$$

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

$$z_1 \geq \dots \geq z_{n^*} > \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta} \geq z_{n^*+1} \geq \dots \geq z_n. \quad (4.20)$$

- else, $f_i^{(\delta)}(z) \stackrel{\text{def}}{=} f_{\pi(i)}^{(\delta)}(z_{\pi(1)}, \dots, z_{\pi(n)})$ where π is any permutation of the players such that $z_{\pi(1)} \geq \dots \geq z_{\pi(n)}$ (i.e., we define $f_i^{(\delta)}$ by extending it symmetrically).

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

Lemma 4.5. $f^{(\delta)}$ is a well-defined allocation function.

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

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

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

Indeed, Eq. (4.21) vacuously holds for $n' = n$. Now take n'' to be the least such index.

Accordingly,

$$\forall i > n'', \quad z_i \leq \frac{\sum_{j=1}^{n''} z_j}{n'' + D_\delta}. \quad (4.21)$$

Next we claim that

$$\forall i \leq n'', \quad z_i > \frac{\sum_{j=1}^{n''} z_j}{n'' + D_\delta}. \quad (4.22)$$

To prove Eq. (4.22), it suffices to consider $i = n''$ because z is non-increasing. Indeed, by the minimality of n'' we know that (“ $n'' - 1$ does not work”, that is) there exists some $j \geq n''$ such that

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

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

At last, combining Eq. (4.21) and Eq. (4.22), and choosing $n^* = n''$, Eq. (4.20) is satisfied.

Next, we prove that n^* is unique. Suppose by way of contradiction that there exist two integers n^\perp and n^\top , with $n^\perp < n^\top$ both satisfying Eq. (4.20). Now define

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

By invoking Eq. (4.20) with n^\top and n^\perp , we deduce that for $i \in \{n^\perp + 1, \dots, n^\top\}$,

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

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

$$\frac{S^\perp}{n^\perp + D_\delta} \geq \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + S^\Delta}{n^\perp + n^\Delta + D_\delta} . \quad (4.23)$$

Let us now show that the second inequality of Eq. (4.23) contradicts the first inequality Eq. (4.23):

$$\begin{aligned} \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + S^\Delta}{n^\perp + n^\Delta + D_\delta} &\Leftrightarrow (n^\perp + n^\Delta + D_\delta)S^\Delta > n^\Delta(S^\perp + S^\Delta) \\ &\Leftrightarrow (n^\perp + D_\delta)S^\Delta > n^\Delta S^\perp \Leftrightarrow \frac{S^\Delta}{n^\Delta} > \frac{S^\perp}{(n^\perp + D_\delta)} . \end{aligned} \quad (4.24)$$

The contradiction establishes the uniqueness of n^* .

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

Assume, again without loss of generality, that $z_1 \geq z_2 \geq \dots \geq z_n$. Eq. (4.20) tells us that $z_i(n^* + D_\delta) - \sum_{j=1}^{n^*} z_j \geq 0$ for each $i \leq n^*$, so (a) follows immediately. As for

(b),

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

□

Lemma 4.6. $f^{(\delta)}$ is monotonic.

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

$$f_n^{(\delta)}(z_{-n} \sqcup z_n^\perp) \leq f_n^{(\delta)}(z_{-n} \sqcup z_n^\top) . \quad (4.25)$$

We will prove Eq. (4.25) in three steps.

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

$$z_n \leq \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \longrightarrow f_n^{(\delta)}(z_{-n} \sqcup z_n) = 0 \text{ (i.e., } n \text{ is a loser)} \quad (4.26)$$

$$z_n > \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta} \longrightarrow f_n^{(\delta)}(z_{-n} \sqcup z_n) > 0 \text{ (i.e., } n \text{ is a winner)} \quad (4.27)$$

To show Eq. (4.26), recall that, in the game with only the first $n - 1$ players bidding z_{-n} , we have n' winners satisfying,

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

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

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

To show Eq. (4.27), we actually prove its contrapositive: namely,

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

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

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

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

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

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

Because of Step 1, we only need to show Eq. (4.25) for z_n^\perp and z_n^\top satisfying $z_n^\top > z_n^\perp > \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta}$. Notice that in such a case, player n is always a winner. Therefore, let $\{1, \dots, n^\perp, n\}$ and $\{1, \dots, n^\top, n\}$ be the winners when the bid profiles are $(z_{-n} \sqcup z_n^\perp)$ and $(z_{-n} \sqcup z_n^\top)$ respectively.

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

$$n^\perp \geq n^\top. \quad (4.28)$$

Assume by way of contradiction that $n^\perp < n^\top$ and. As in Lemma 4.5, set $n^\Delta \stackrel{\text{def}}{=} n^\top - n^\perp$, $S^\perp \stackrel{\text{def}}{=} \sum_{j=1}^{n^\perp} z_j$ and $S^\top \stackrel{\text{def}}{=} \sum_{j=1}^{n^\top} z_j = S^\perp + S^\Delta$. Then each player i , with $n^\perp \leq i < n^\top$, is a loser when the bid profile is $(z_{-n} \sqcup z_n^\perp)$ while a winner when the bid profile is $(z_{-n} \sqcup z_n^\top)$; in particular,

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

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

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

but this is already a contradiction, since the right hand side is equivalent to

(using a similar technique as Eq. (4.24)):

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

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

We now use the fact that $n^\perp \geq n^\top$ to obtain Eq. (4.25) for such z_n^\perp and z_n^\top satisfying $z_n^\top > z_n^\perp > \frac{\sum_{j=1}^{n'} z_j}{n' + D_\delta}$.

• **Step 3.** We now prove Eq. (4.25).

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

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

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

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} = \frac{S^\top + S^\Delta + z_n^\perp}{n^\top + n^\Delta + 1 + D_\delta}. \quad (4.29)$$

But this is equivalent to (again using the same technique as Eq. (4.24))

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\top + z_n^\perp}{n^\top + 1 + D_\delta}. \quad (4.30)$$

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

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

$$\begin{aligned}
&= C_2 \cdot \left(S^\top(n^\top + 1 + D_\delta)(z_n^\top - z_n^\perp) + S^\Delta(n^\top + 1 + D_\delta)z_n^\top - n^\Delta(S^\top + z_n^\top)z_n^\perp \right) \\
&\geq C_2 \cdot \left(S^\top(n^\top + 1 + D_\delta)(z_n^\top - z_n^\perp) + S^\Delta(n^\top + 1 + D_\delta)z_n^\top - n^\Delta(S^\top + z_n^\perp)z_n^\top \right) \geq 0
\end{aligned}$$

Here the last inequality has used $z_n^\top - z_n^\perp \geq 0$ and $S^\Delta(n^\top + 1 + D_\delta) - n^\Delta(S^\top + z_n^\perp) > 0$ (by Eq. (4.30)). \square

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

Lemma 4.7. $f^{(\delta)}$ is 1-distinguishably monotonic.

Proof. We already know from Lemma 4.6 that $f^{(\delta)}$ is monotonic. Also, the integrability of $f^{(\delta)}$ is obvious, because $f^{(\delta)}$ is piecewise continuous, and there are at most n pieces, as the number of winners decreases when z_n increases (recall Eq. (4.28)). We are therefore left to prove the “distinguishability condition”.

Fix a player $i \in [n]$ and two distinct valuations $v_i, v'_i \in \{0, 1, \dots, B\}$, and assume that $v_i < v'_i$. Define $v_{-i} \stackrel{\text{def}}{=} (v_i, v_i, \dots, v_i)$, then:

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

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

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

as desired. \square

Lemma 4.8. $f^{(\delta)}$ is δ -good.

Proof. We already know from Lemma 4.7 that $f^{(\delta)}$ is 1-DM. Therefore, in order to prove that $f^{(\delta)}$ is δ -good, we only need to show that Eq. (4.14) holds. We will actually prove that Eq. (4.14) holds not only for the discrete cube $\{0, 1, \dots, B\}^{[n]}$ but also in the continuous cube $[0, B]^{[n]}$.

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

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

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

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

satisfying Eq. (4.14), where the last inequality is due to $k > n^*$ and Eq. (4.20).

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

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

again satisfying Eq. (4.14). □

4.5.3 Our Mechanism $M_{\text{opt}}^{(\delta)}$

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

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

Proof. By Lemma 4.8, the function $f^{(\delta)}$ from Definition 4.4 is a (well-defined) allocation function that is also δ -good. Therefore, invoking Lemma 4.3, the mechanism $M_{\text{opt}}^{(\delta)} \stackrel{\text{def}}{=} M_{f^{(\delta)}}$ yields the target social welfare. ■

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

Claim 4.9. *The outcome function F of $M_{\text{opt}}^{(\delta)}$ is efficiently computable.*

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

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

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

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

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

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

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

This is again not hard, by using a simple line sweep method. One can start from $v_n = \frac{\sum_{j=1}^{n'} v_j}{n'+D_\delta}$ and move v_n upwards. At any moment, one can calculate the earliest time that Eq. (4.20) is violated, and claim that another piece of $f_n^{(\delta)}$ is found. \square

Chapter 5

Multi-Good Auctions

We now raise the bar, and consider mechanism design in the approximate-valuation world for combinatorial auctions. We start with its definitional difference from that for single-good auctions.

5.1 The Model

As before, we start with the auction contexts in (*truly-*)*combinatorial auctions*. For $\delta \in [0, 1)$, a δ -*approximate (auction) context* consists of the following components.

- $[n] = \{1, 2, \dots, n\}$, the set of *players*.
- $[m] = \{1, \dots, m\}$, the set of *goods*.
- θ , the *true-valuation profile*, where each $\theta_i: 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$ with $\theta_i(\emptyset) = 0$.
- $\Omega = \mathcal{A}_n \times \mathbb{R}^{[m]}$, where \mathcal{A}_n is the set of profiles $A = (A_0, A_1, \dots, A_n)$ that are partitions of $[m]$, is the set of *outcomes*; if $(A, P) \in \Omega$ is an outcome, then we refer to A as its *allocation* (i.e., A_i is the subset of the goods assigned to i , and A_0 is the set of unallocated goods) and to P as its profile of *prices*.
- U , the profile of *utility functions*; each U_i maps any outcome (A, P) to $\theta_i(A_i) - P_i$.
- K , the *approximate-valuation profile*, where, for all i , $K_i: 2^{[m]} \rightarrow 2^{\mathbb{R}_{\geq 0}}$ gives the list of possible valuations of player i for each non-empty subset of the goods (i.e., it always satisfies $K_i(\emptyset) = \{0\}$). Each K_i is δ -*approximate*: for all $S \subseteq [m]$

and $S \neq \emptyset$, we have

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

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

Notice that C is fully specified by n, m, δ, θ, K , that is $C = (n, m, \delta, \theta, K)$. Similarly as before, in a context $C = (n, m, \delta, \theta, K)$ each player i only knows K_i and that $\theta_i \in K_i$ (which means $\theta_i(S) \in K_i(S)$ for all S), but not necessarily δ or K_j for other player $j \neq i$.

We remark here that slightly different from the definitions in single-good auctions, we let players’ set of possible valuations be $\mathbb{R}_{\geq 0}$ for each subset of the goods, rather than a discrete set $\{0, 1, \dots, B\}$. This greatly simplifies the proof of one of our theorems.

The *social welfare* (function) SW is defined as $SW(\theta, (A, P)) \stackrel{\text{def}}{=} \sum_{i \in [n]} \theta_i(A_i)$ for every true-valuation profile θ and outcome (A, P) . The *maximum* social welfare of a true-valuation profile θ , $MSW(\theta)$, is defined to be the maximum of $SW(\theta, (A, P))$ over all possible outcomes $(A, P) \in \Omega$.

5.2 Our Results

5.2.1 A First Impossibility Result

In dominant-strategies, our impossibility result Theorem 1 for single-good auctions dramatically generalizes to truly combinatorial auctions.

Theorem 4. *For all $n, \delta \in (0, 1)$, and (possibly probabilistic) very-weakly-dominant-strategy-truthful mechanism $M = (S, F)$, there exists a δ -approximate-valuation profile K and a true-valuation profile $\theta \in K$ such that*

$$\mathbb{E} \left[SW(\theta, F(K)) \right] \leq \frac{1}{n2^{m-1}} MSW(\theta) .$$

(The proof of Theorem 4 can be found in Section 6.2.)

Further notice that a fraction $\frac{1}{n(2^m-1)}$ of the maximum social welfare can be trivially achieved by the “stupid” very-weakly-dominant-strategy mechanism that, ignoring all bids, assigns a random non-empty subset of the goods to a random player, and leaves all other goods unassigned. Our theorem is thus asymptotically tight (up to a factor of 2).

Exactly for the same reason as Theorem 1, Theorem 4 is *non-trivial* because each player i is allowed to for instance bid his own K_i , then in principle there might be a smart way for a dominant-strategy mechanism to turn such “truthful” bids into an outcome of reasonable welfare!

5.2.2 A Second Impossibility Result

Recall that for combinatorial auctions in the exact-valuation world, the VCG mechanism guarantees perfect social welfare in dominant strategies. The same is true for the same mechanism but in undominated strategies. That is,

Fact 5. *In the exact-valuation world, the VCG mechanism guarantees maximum social welfare in undominated strategies.*

(We have been unable to ascertain whether this fact was previously known, but we are eager to attribute it properly. Meanwhile, to avoid any doubts, we provide a proof of it in Section 6.1.)

The above fact legitimizes asking whether our second, positive result for single-good auctions (i.e., result Theorem 2) generalizes. That is: can the VCG mechanism, in approximate-valuation combinatorial auctions, guarantee some reasonable efficiency in undominated strategies? Our answer is a resounding NO.

Theorem 6. *Let $M = (S, F)$ be a VCG mechanism with any tie-breaking rule. Then, for all $n \geq 2$, $m \geq 2$ and $\delta \in (0, 1)$, there exist a δ -approximate-valuation profile K , a true-valuation profile $\theta \in K$, and a strategy profile $s \in \text{UDed}(K)$ such that:*

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

(The proof of this theorem is very technical, and can be found in Section 6.3.)¹

In other words, while the VCG essentially is the top mechanism for achieving efficiency in traditional combinatorial auctions (even in undominated strategies, as we pointed out in Fact 5), in the approximate-valuation world it is close to the bottom. In fact, it is outperformed by quite simple-minded undominated-strategy mechanisms. For example, the mechanism that first asks each player to report a *single* set along with a value for it, and then runs the second-price mechanism, guarantees a $\frac{1}{m} \left(\frac{1-\delta}{1+\delta}\right)^2$ -fraction of the maximum social welfare in undominated strategies.

One might certainly conceive more sophisticated uses of “dimensionality-reduction techniques” to improve the performance of the above simple-minded (i.e., “dimension-1”) mechanism, but we do not believe that such approaches can yield a social welfare guarantee that solely depends on the approximation factor of the players’ internal knowledge (i.e., solely on δ). Actually, the author is aware of a technique to extend this impossibility result at least to all Maximal-in-Range (MIR) mechanisms.

¹Let us note that the exponentially poor performance of the VCG continues to hold even if the players are μ -utility-indifferent for $\mu < \delta$. (If $\mu > \delta$, then the performance of the VCG is restored, but then we will be in a “degenerate” approximate-valuation setting, since a player’s uncertainty is “drowned” by his utility indifference.)

Chapter 6

Multi-Good Auctions:

Missing Proofs

6.1 Proof for Fact 5

In the exact-valuation world, it is well known that the VCG mechanism guarantees full social welfare in (very-weakly) dominant strategies: if every player reports the truth then full social welfare is guaranteed and, moreover, reporting the truth is a very-weakly dominant strategy.

However, when analyzing the same mechanism but in undominated strategies, we do not assume that players will report the truth just because it is undominated. For instance, when there are two goods $[m] = \{1, 2\}$, a player with true valuation ($\{1\} = 6, \{2\} = 10, \{1, 2\} = 8$) may also consider to bid ($\{1\} = 6, \{2\} = 10, \{1, 2\} = 7$), since underbidding on $\{1, 2\}$ does not affect the allocation in the VCG mechanism. Indeed, as we will formalize later in Claim 6.1, when players deviate to some undominated non-truthful strategy, the allocation “does not change”. We will then use Claim 6.1 to prove Fact 5. Details now follow:

Fact 5 (restated). *In the exact-valuation setting, the VCG mechanism with randomized tie breaking rule¹ guarantees full social welfare also in undominated strategies.*

¹See Remark 6.2 for a brief discussion about how tie breaking rules affect our statement.

Recall that in the VCG mechanism, player i 's set of possible strategies S_i equals to his set of possible valuations $\Theta_i := \mathbb{R}_{\geq 0}^{2^{|m|}}$. In this paper we consider truly combinatorial auctions so that each player can have an arbitrary non-negative valuation function. Often in the literature, people have also studied restricted cases of Θ_i (e.g., for instance letting Θ_i contain only submodular or matroid-rank sum functions); we emphasize that our proof below also works for *arbitrary* choices of Θ_i 's as a subset of $\mathbb{R}_{\geq 0}^{2^{|m|}}$.

Let the true valuation profile of the players be $\theta \in \Theta_i$. It is straightforward to see that for each player i bidding his truth is an undominated strategy $\theta_i \in \text{UDed}_i(\theta_i)$, since a very-weakly-dominant strategy is always not-weakly-dominated. By Lemma 7.9, this implies that $\text{UDed}_i(\theta_i) = \text{Dnt}_i(\theta_i)$ so each undominated strategy is also a (very-weakly-)dominant one for each player i .

Let $A(v)$ be the set of allocations that maximizes the social welfare according to bidding profile v , i.e.,

$$A(v) = \arg \max_{A=(A_0, A_1, \dots, A_n)} \left\{ \sum_i v_i(A_i) \right\} .$$

The VCG mechanism with *randomized tie breaking rule* will output a random one among the finite number of choices in $A(v)$, each with some positive probability.

The proof of Fact 5 now relies on the following crucial observation.

Claim 6.1. *In the VCG mechanism with randomized tie breaking rule, for any player i with true valuation θ_i and some other undominated strategy $\theta'_i \in \text{UDed}_i(\theta_i)$, we have*

$$\forall v_{-i} \in \Theta_{-i}, \quad A(\theta'_i \sqcup v_{-i}) \subseteq A(\theta_i \sqcup v_{-i}) . \quad (6.1)$$

Proof. Fix some v_{-i} and let $\{(p_a, a)\}_{a \in A(\theta'_i \sqcup v_{-i})}$ be the probability distribution over allocations that the VCG mechanism outputs on input bid profile $\theta'_i \sqcup v_{-i}$, and similarly define $\{(q_b, b)\}_{b \in A(\theta_i \sqcup v_{-i})}$ for that on input bid profile $\theta_i \sqcup v_{-i}$.

Since $\theta'_i \in \text{UDed}_i(\theta_i) = \text{Dnt}_i(\theta_i)$, θ'_i very-weakly dominates θ_i , i.e.,

$$\mathbb{E}U(\theta_i, \text{VCG}(\theta'_i \sqcup v_{-i})) \geq \mathbb{E}U(\theta_i, \text{VCG}(\theta_i \sqcup v_{-i})) ,$$

and by writing down the utilities and canceling out common terms, we have

$$\sum_{a \in A(\theta'_i \sqcup v_{-i})} p_a(\theta_i(a) + v_{-i}(a)) \geq \sum_{b \in A(\theta_i \sqcup v_{-i})} q_b(\theta_i(b) + v_{-i}(b)) .$$

If we use OPT to denote the maximum social welfare on input bid profile $\theta_i \sqcup v_{-i}$, then obviously the right hand side equals to OPT, as all b 's are by definition maximizers to this welfare. On the contrary, for every allocation a on the left, we have: $\theta_i(a) + v_{-i}(a) \leq \text{OPT}$ by the definition of OPT. This, along with the fact that $p_a > 0$ for all a , gives us the conclusion that all allocation a on the left will satisfy $\theta_i(a) + v_{-i}(a) = \text{OPT}$, and thus $a \in A(\theta_i \sqcup v_{-i})$. This concludes that Eq. (6.1) holds. \square

Now we come back to the proof of Fact 5. Recall that $\theta = (\theta_1, \dots, \theta_n)$ is the true valuation profile of the players, and for each $i \in [n]$, let $\theta'_i \in \text{UDed}_i(\theta_i)$ be some arbitrary undominated strategy to be played by player i . By repeatedly applying the above claim, we have

$$A(\theta'_1 \sqcup \theta'_2 \sqcup \dots \sqcup \theta'_n) \subseteq A(\theta_1 \sqcup \theta'_2 \sqcup \dots \sqcup \theta'_n) \subseteq \dots \subseteq A(\theta_1 \sqcup \theta_2 \sqcup \dots \sqcup \theta_n) .$$

This says, when players are playing undominated strategies, the possible allocations that the VCG mechanism may output, is still a subset of all social welfare maximizers with respect to the true valuations. This completes the proof that the VCG mechanism with randomized tie breaking rule guarantees full social welfare when players play arbitrary undominated strategies. \blacksquare

Remark 6.2. We now make several remarks:

- Fact 5 holds for arbitrary valuation space Θ_i , including for instance a discrete space $\Theta_i = \{0, 1, 2, \dots, B\}^{2^{|m|}}$, or a submodular space. It can also be extended to maximal-in-range (MIR) or maximal-in-distributional-range (MIDR) mechanisms. As a consequence, noticing that undominated-strategy implementation is provably a safer solution concept than dominant-strategy one, Fact 5 immediately strengthens all prior results on designing VCG-type mechanisms; see for instance [13].

- If the VCG mechanism breaks ties deterministically, or probabilistically but with a probability zero at some of the social welfare maximizers, then our above proof no longer holds. However, Fact 5 still holds when Θ_i is continuous. The proof is different and much more involved. We omit it in this thesis.
- If the VCG mechanism breaks ties deterministically, and at the same time $\Theta_i = \{0, 1, 2, \dots, B\}^{2^{[m]}}$ is discrete, then the maximum social welfare might not be guaranteed. As a simple example in a single-good auction, if the VCG mechanism (i.e., the Vickery auction) breaks tie lexicographically, the first player may underbid by an additive 1, and the last player may overbid by an additive 1, resulting in a maximum additive 1 loss in the social welfare.

6.2 Proof for Theorem 5

We present a slightly stronger version of the theorem, which assumes that each player’s valuations on each non-empty subset of the goods can only reside in a discrete set $\{0, 1, \dots, B\}$, instead of the entire $\mathbb{R}_{\geq 0}$. It is easy to see that by letting $B \rightarrow \infty$, the same theorem also holds for $\mathbb{R}_{\geq 0}$.

The main idea of the proof is similar to the one that we used in the case of single-good auctions in Theorem 1. What lets us achieve a stronger negative result in the multi-good case is the fact that now we have many more valuations (namely, $n(2^m - 1)$ in total) among which the probability mass is spread.

Theorem 4 (restated for discrete valuations). *For all $n, m, \delta \in (0, 1)$, $B > \frac{3-\delta}{2\delta}$, and (possibly probabilistic) very-weakly-dominant-strategy-truthful mechanisms $M = (S, F)$, there exist a δ -approximate-valuation profile K and a true-valuation profile $\theta \in K$ such that*

$$\mathbb{E}[\text{SW}(\theta, F(K))] \leq \left(\frac{1}{n2^{m-1}} + \frac{\lfloor \frac{3-\delta}{2\delta} \rfloor + 1}{B} \right) \text{MSW}(\theta) .$$

Proof. We want to prove a similar auxiliary claim as Claim 4.1, that says if a player reports a δ -approximate valuation that is “sufficiently high”, his “winning probabili-

ty” and expected price remain constant. We need to introduce several new definitions before we can even state our claim in this complicated combinatorial scenario.

For every δ -approximate-valuation profile \tilde{K} and possible outcome $\omega \in \Omega$, define $F_\omega^A(\tilde{K})$ to be the probability that mechanism $F(\tilde{K})$ chooses the outcome ω . For every player $i \in [n]$, subset $S \subseteq [m]$, and outcome $\omega \in \Omega$, we say that the pair (i, S) is *consistent* with ω , denoted $(i, S) \sim \omega$, if $\omega = (A, P)$ and $A_i = S$. Next, for every δ -approximate-valuation profile \tilde{K} , player $i \in [n]$, and subset $S \subseteq [m]$, we define $F_{i,S}^A(\tilde{K})$ to be the probability that player i receives subset S , i.e.,

$$F_{i,S}^A(\tilde{K}) \stackrel{\text{def}}{=} \sum_{\substack{\omega \in \Omega \\ (i,S) \sim \omega}} F_\omega^A(\tilde{K}) .$$

This last definition is motivated by the fact that, in a combinatorial auction, player i is only interested in his own allocation A_i , and is indifferent to A_{-i} . We now show the variant of Claim 4.1 as follows:

Claim 6.3. *For all player i , integer $x \in (\frac{3-\delta}{2\delta}, B]$, δ -approximate-valuation profile \tilde{K} , δ -approximate valuation \tilde{K}'_i , and non-empty $S \subseteq [m]$. If $\tilde{K}_i(S) = \delta[x]$, $\tilde{K}'_i(S) = \delta[x+1]$, while $\tilde{K}_i(T) = \tilde{K}'_i(T)$ for all $T \neq S$, then*

$$\begin{aligned} F_{i,S}^A(\tilde{K}_i \sqcup \tilde{K}_{-i}) &= F_{i,S}^A(\tilde{K}'_i \sqcup \tilde{K}_{-i}) \text{ and} \\ F_i^P(\tilde{K}_i \sqcup \tilde{K}_{-i}) &= F_i^P(\tilde{K}'_i \sqcup \tilde{K}_{-i}) . \end{aligned}$$

Proof. We consider two cases:

- If player i has approximate type \tilde{K}_i then reporting \tilde{K}_i very-weakly dominates reporting \tilde{K}'_i :

$$\begin{aligned} \forall \theta_i \in \tilde{K}_i : \quad & \sum_{T \subseteq [m]} \left(F_{i,T}^A(\tilde{K}_i \sqcup \tilde{K}_{-i}) \cdot \theta_i(T) \right) - F_i^P(\tilde{K}_i \sqcup \tilde{K}_{-i}) \\ & \geq \sum_{T \subseteq [m]} \left(F_{i,T}^A(\tilde{K}'_i \sqcup \tilde{K}_{-i}) \cdot \theta_i(T) \right) - F_i^P(\tilde{K}'_i \sqcup \tilde{K}_{-i}) . \end{aligned}$$

- If player i has approximate type \tilde{K}'_i then reporting \tilde{K}'_i very-weakly dominates

reporting \tilde{K}_i :

$$\begin{aligned} \forall \theta'_i \in \tilde{K}'_i: & \sum_{T \subseteq [m]} \left(F_{i,T}^A(\tilde{K}'_i \sqcup \tilde{K}_{-i}) \cdot \theta'_i(T) \right) - F_i^P(\tilde{K}'_i \sqcup \tilde{K}_{-i}) \\ & \geq \sum_{T \subseteq [m]} \left(F_i^A(\tilde{K}_i \sqcup \tilde{K}_{-i}) \cdot \theta'_i(T) \right) - F_i^P(\tilde{K}_i \sqcup \tilde{K}_{-i}) . \end{aligned}$$

On one hand, we can choose $\theta_i(S) = x$ and $\theta'_i(S) = x + 1$, and for all $T \neq S$ choose $\theta_i(T) = \theta'_i(T)$ to be some arbitrary point in $\tilde{K}_i(T) = \tilde{K}'_i(T)$. Summing up the two inequalities, we get:

$$F_{i,S}^A(\tilde{K}'_i \sqcup \tilde{K}_{-i}) \geq F_{i,S}^A(\tilde{K}_i \sqcup \tilde{K}_{-i}) .$$

On the other hand, we can choose $\theta_i(S) = \lfloor x(1 + \delta) \rfloor \in \tilde{K}_i(S)$ and $\theta'_i(S) = \lceil (x + 1)(1 - \delta) \rceil \in \tilde{K}'_i(S)$, and for all $T \neq S$ choose $\theta_i(T) = \theta'_i(T)$ to be some arbitrary point in $\tilde{K}_i(T) = \tilde{K}'_i(T)$.² Again summing the two inequalities, most of the terms cancel, and we are left with the following:

$$\left(F_{i,S}^A(\tilde{K}_i \sqcup \tilde{K}_{-i}) - F_{i,S}^A(\tilde{K}'_i \sqcup \tilde{K}_{-i}) \right) \cdot (\lfloor x(1 + \delta) \rfloor - \lceil (x + 1)(1 - \delta) \rceil) \geq 0 .$$

Therefore, whenever $x > \frac{3-\delta}{2\delta}$, we always have $\lfloor x(1 + \delta) \rfloor - \lceil (x + 1)(1 - \delta) \rceil > 0$ and thus $F_{i,S}^A(\tilde{K}_{i,S} \sqcup \tilde{K}_{-i}) = F_{i,S}^A(\tilde{K}'_i \sqcup \tilde{K}_{-i})$. Finally, going back to the two inequalities for very-weak dominance, we can also deduce that $F_i^P(\tilde{K}_i \sqcup \tilde{K}_{-i}) = F_i^P(\tilde{K}'_i \sqcup \tilde{K}_{-i})$, as desired. \square

We can now go back to the proof of the theorem. Define $c \stackrel{\text{def}}{=} \lfloor \frac{3-\delta}{2\delta} \rfloor + 1$ and $\hat{K}_i(T) \stackrel{\text{def}}{=} \delta[c]$ for all nonempty $T \subseteq [m]$ and players $i \in [n]$. Because the mechanism assigns disjoint outcomes with a maximum total probability of 1, we have:

$$\sum_{i \in [n]} \sum_{\substack{T \subseteq [m] \\ 1 \in T}} F_{i,T}^A(\hat{K}) \leq 1 .$$

Again, the events in the summation are disjoint, because there is only one good $1 \in [m]$ and it can be assigned to only one of the n players. Also because $|\{T \subseteq G : 1 \in T\}| = 2^{m-1}$, at least one of the probabilities, say $F_{1,S}^A(\hat{K})$, is at most $\frac{1}{n2^{m-1}}$.

²Again, as long as $x > \frac{1}{2\delta}$, it is guaranteed that, for these choices, $\theta_i \in \tilde{K}_i$ and $\theta'_i \in \tilde{K}'_i$. But later we will choose $x > \frac{3-\delta}{2\delta}$, so we are safe.

Now define $K_1(S) \stackrel{\text{def}}{=} \delta[B]$ and $K'_1(T) \stackrel{\text{def}}{=} \delta[c]$ for all nonempty $T \neq S$, and $K_i(T) = \delta[c]$ for all players $i \neq 1$ and nonempty $T \subseteq [m]$. Invoking Claim 6.3 multiple times with player 1 and the subset S , K_{-i} of this proof, and x going from c to B , we obtain that

$$F_{1,S}^A(K) = F_{1,S}^A(\widehat{K}) \leq \frac{1}{m2^{m-1}} .$$

Now suppose that the true approximate type profile of the players is K . Then, for the choice of the true type profile $\theta = (\theta_1, \dots, \theta_n)$ with $\theta_1(S) = B$ and $\theta_1(T) = c$ for all nonempty $T \neq S$, and $\theta_i(T) = c$ for all players $i \neq 1$ and nonempty $T \subseteq [m]$, we get the following social welfare:

$$\begin{aligned} \mathbb{E}[\text{SW}(\theta, F(K))] &\leq \frac{B}{n2^{m-1}} + \left(1 - \frac{1}{n2^{m-1}}\right) \cdot c \leq \left(\frac{1}{n2^{m-1}} + \frac{c}{B}\right) \cdot B \\ &\leq \left(\frac{1}{n2^{m-1}} + \frac{c}{B}\right) \cdot \text{MSW}(\theta) . \end{aligned}$$

■

6.3 Proof for Theorem 6

In this section we prove the following theorem:

Theorem 6. *Let $M = (S, F)$ be a VCG mechanism with any tie-breaking rule. Then, for all $n \geq 2$, $m \geq 2$ and $\delta \in (0, 1)$, there exist a δ -approximate-valuation profile K , a true-valuation profile $\theta \in K$, and a strategy profile $s \in \text{UDed}(K)$ such that:*

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

6.3.1 Proof Sketch

The very high level reason behind Theorem 6 is that the “argmax function” of the VCG mechanism, which selects the best allocation of the goods given the reported valuations, is extremely “unstable” in high dimensions (i.e., when the number of goods is large) with δ -approximate valuations.

Translating the above idea into a proof is not easy, but let us try to give its flavor.

“Highly-deviating” undominated strategies. Which strategies are undominated in the VCG mechanism? Consider a player i such that, for a fixed value x , $K_i(S) = \delta[x] = [(1 - \delta)x, (1 + \delta)x]$ for all non-empty subset S of the goods (i.e., his uncertainty is δ -clustered around x everywhere). Now let us consider two possible strategies of player i .

The first consists of reporting the valuation w_i such that $w_i(S) = x(1 + 10\delta)$ for every non-empty S . This strategy is intuitively dominated, since it deviates too much from the “approximate truth” K_i . Indeed, it can be verified that w_i is dominated by a strategy of reporting $x(1 + \delta)$ for every subset S .

Let us now consider a second example. Call a permutation of all non-empty subsets of the goods $\pi = (\pi_1, \dots, \pi_{2^m-1})$ *proper* if $j < k \Rightarrow \pi_j \not\supseteq \pi_k$.³ Then, for any proper π , define the valuation v_i as follows:

$$\forall j \in \{1, 2, \dots, 2^m - 1\}, \quad v_i(\pi_j) \stackrel{\text{def}}{=} x + 2(j - 1)\delta x . \quad (6.2)$$

Surprisingly, this “high deviating” strategy is actually *undominated*.

Let us now give a very rough explanation to why this is so. Consider first a pure strategy $v'_i \neq v_i$. Then, there must exist some consecutive sets π_j and π_{j+1} , such that $v_i(\pi_j) - v_i(\pi_{j+1}) > v'_i(\pi_j) - v'_i(\pi_{j+1})$. Assume by way of contradiction now that v_i is dominated by v'_i . If the rest of the players happen to bid such that only subsets π_j or π_{j+1} could possibly be given to player i by the VCG mechanism, then only the *difference* between $K_i(\pi_j)$ and $K_i(\pi_{j+1})$ matters in terms of player i 's utility. However, we see in Eq. (6.2) that the relative difference between $v_i(\pi_j)$ and $v_i(\pi_{j+1})$ is not even more than $2\delta x$, resulting in a contradiction. A very careful analysis will justify that v_i cannot in fact be dominated, by *any* mixed strategy.

More generally (and loosely speaking), for each proper permutation π , any bid v_i satisfying the following constraints is undominated:

$$\forall 1 \leq i \leq 2^m - 1, \quad v_i(\pi_j) - v_i(\pi_{j+1}) \geq \min K_i(\pi_j) - \max K_i(\pi_{j+1}) ,$$

which is a full-dimensional triangular cylinder. Hence, by taking the union over

³When $m = 3$ 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 proper permutations.

all proper permutations (of which there are exponentially many), we obtain a full-dimensional subset of the bid space, each of whose points is an undominated strategy of player i ; we call it a “bird” and each triangular cylinder a wing of the bird. For more details, see the Bird Lemma 6.4.

(A characterization. In fact, the reverse is also true: namely, if there does not exist a wing containing a given bid v_i , then v_i must be dominated. This gives the first *characterization* of the undominated strategies of the VCG in an approximate-valuation setting. We omit proving this direction in this thesis, because it is not necessary for the proof of Theorem 6.)

A hard instance. Finally, in order to show that the social welfare guaranteed by the VCG mechanism is not more than $(\frac{1-\delta}{1+\delta})^{2^m-2}$, as claimed in Theorem 6, we design two “highly-deviating” undominated strategies for two different players, one strategy significantly overbidding and one strategy significantly underbidding, in order to “confuse” the VGC mechanism. This hard instance will be chosen in Section 6.3.4.

6.3.2 The Bird Lemma

Lemma 6.4 (Bird Lemma). *In the VCG mechanism, no matter how the ties are broken, for a player, if his approximate valuation K and his bid v satisfy the following requirements, then $v \in \text{UDed}(K)$ for him.*

1. *Each $K(S)$ is an interval: $\forall S \subseteq [m]$ and $S \neq \emptyset$, $K(S) = [K(S)^\perp, K(S)^\top]$ for some real $K(S)^\perp, K(S)^\top$.*
2. *K^\perp and K^\top are weakly monotone:*

$$\forall \emptyset \neq S \subseteq T \subseteq [m], \quad K(S)^\perp \leq K(T)^\perp ,$$

$$\forall \emptyset \neq S \subseteq T \subseteq [m], \quad K(S)^\top \leq K(T)^\top .$$

3. *v is strictly monotone:*

$$\forall \emptyset \neq S \subseteq T \subseteq [m], \quad v(S) < v(T) .$$

4. At least one coordinate of v is above (resp. below) its lower (resp. upper) bound:

$$\exists S' \in [m], \quad v(S') \leq K(S')^\top \quad (6.3)$$

$$\exists S'' \in [m], \quad v(S'') \geq K(S'')^\perp \quad (6.4)$$

5. There exists a proper permutation π of all non-empty subsets of $[m]$, such that:

$$\forall 1 \leq i \leq 2^m - 1, \quad v(\pi_i) - v(\pi_{i+1}) \geq K(\pi_i)^\perp - K(\pi_{i+1})^\top . \quad (6.5)$$

Here we assume that $\pi(2^m) = \pi(1)$.

6.3.3 Proof of the Bird Lemma

Let the player that we consider be the first player. Assume by contradiction that his bidding strategy v is weakly dominated by some mixed strategy $\{p_j, v^{(j)}\}_j$, where the probabilities p_j 's sum up to one and $v^{(j)} \neq v$ for all j . Our goal is to provide a witness bid of the other player w^4 , and a witness true valuation $\theta \in K$ such that, if U is the utility function for the first player, then

$$U(\theta, \text{VCG}(v \sqcup w)) > \sum_j p_j U(\theta, \text{VCG}(v^{(j)} \sqcup w)) . \quad (6.6)$$

This will contradict the fact that the mixed strategy $\{p_j, v^{(j)}\}_j$ weakly dominates v . We shall, however, create such witness $w : 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$ and $\theta \in K$ based on different cases.

Notation.

- We call the player bidding v the first player, and the player bidding w the second (or the witness) player. We drop the subscript for notational simplicity.
- We say that the allocation of $\text{VCG}(v \sqcup w)$ is (S, T) if the first player receives S and the second player receives T .
- We use $\text{SW}[(S, T)] = v(S) + w(T)$ to denote the “social welfare” of the allocation, as if both players have exactly the reported bids $v \sqcup w$ as their true valuations.

⁴Notice that it suffices to consider only a two-player game, since if we have more than two players, we can let the rest of them bid 0 as a witness.

- Since VCG mechanism is maximizing social welfare, we have that $\text{SW}[\text{VCG}(v \sqcup w)] = \max_{(S,T)} \{v(S) + w(T)\}$ where the maximization is over all partition (S, T) of the good set $[m]$.
- For notational simplicity, given a bid v , we define its monotonizer \tilde{v} to be such that:

$$\tilde{v}(S) \stackrel{\text{def}}{=} \max_{T \subseteq S} v(T) .$$

Now, among the following inequalities, at least one of them cannot hold:

$$\begin{cases} v(\pi_{i+1}) - v(\pi_i) \leq \min_j \left\{ \tilde{v}^{(j)}(\pi_{i+1}) - \tilde{v}^{(j)}(\pi_i) \right\}, & \forall 1 \leq i \leq 2^m - 1 \\ v(S') \leq \min_j \left\{ \tilde{v}^{(j)}(S') \right\} \\ v(S'') \geq \max_j \left\{ \tilde{v}^{(j)}(S'') \right\} \end{cases} \quad (6.7)$$

We now show that if all inequalities above hold, there must be a contradiction.

Indeed, we can infer from the first inequality that for each i and j , $v(\pi_{i+1}) - v(\pi_i) \leq \tilde{v}^{(j)}(\pi_{i+1}) - \tilde{v}^{(j)}(\pi_i)$, but this sum up to $0 \leq 0$ for all $1 \leq i \leq 2^m - 1$. This means, all such inequalities must be tight, so for each j , $v^{(j)}$ must be the same as v only up to a constant shift. In other words,

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

But substituting this into the second and third inequality in Eq. (6.7), we know that $0 \leq \min_j c^{(j)}$ and $0 \geq \max_j c^{(j)}$, and therefore $c^{(j)}$'s must all be 0, contradicting the fact that $v^{(j)} \neq v$.

Therefore, one of the three kinds of inequalities in Eq. (6.7) must fail, and according to which one of them fails, we have three subcases. Now, we will show that for each possible case, Eq. (6.6) holds, and therefore the bidding strategy v cannot be weakly dominated.

Case 1

Suppose that the first inequality of Eq. (6.7) is broken. Without loss of generality, we assume that it is broken for $i = 1$:

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

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

$$v(\pi_2) - v(\pi_1) > \Delta > \widetilde{v}^{(j^*)}(\pi_2) - \widetilde{v}^{(j^*)}(\pi_1) , \quad (6.8)$$

and for every $j \notin J$:

$$\widetilde{v}^{(j)}(\pi_2) - \widetilde{v}^{(j)}(\pi_1) > \Delta . \quad (6.9)$$

Now, we set the witness bid 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 $\overline{\pi}_1 = \emptyset$ or $\overline{\pi}_2 = \emptyset$ later, since we cannot set the second player to have non-zero valuation on an empty set. We claim that:

Claim 6.5. *If $\overline{\pi}_1 \neq \emptyset$ and $\overline{\pi}_2 \neq \emptyset$:*

- a. *The allocation of VCG($v \sqcup w$) is $\omega = (\pi_2, \overline{\pi}_2)$.*
- b. *For all $j^* \in J$, the allocation of VCG($v^{(j^*)} \sqcup w$) 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 VCG($v^{(j)} \sqcup w$) 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 bids w , if $T \notin \{\overline{\pi}_1, \overline{\pi}_2\}$, then $\text{SW}[(S, T)]$ does not contain the big term H and is thus smaller than any $\text{SW}[\omega]$ 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, $\text{SW}[\omega] = v(\pi_2) + H$. 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 $v(\pi_1) + w(\overline{\pi}_1) =$

$v(\pi_1) + H + \Delta < v(\pi_2) + H = \text{SW}[\omega]$, using Eq. (6.8). In sum, $\omega = (\pi_2, \bar{\pi}_2)$ must be the allocation of the VCG mechanism.

- b. In this case, $\text{SW}[\omega] = \widetilde{v^{(j^*)}}(\pi_1) + H + \Delta$. For the allocation of $(S, \bar{\pi}_1)$, 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 $(S, \bar{\pi}_2)$, $\text{SW}[(S, \bar{\pi}_2)] \leq \widetilde{v^{(j^*)}}(\pi_2) + H < \widetilde{v^{(j^*)}}(\pi_1) + H + \Delta = \text{SW}[\omega]$ (using Eq. (6.8)) is worse than the choice of ω . So the allocation must be of the desired form.
- c. In this case, $\text{SW}[\omega] = \widetilde{v^{(j)}}(\pi_2) + H$. For the allocation of $(S, \bar{\pi}_1)$, we have that $\text{SW}[(S, \bar{\pi}_1)] \leq \widetilde{v^{(j)}}(\pi_1) + H + \Delta < \widetilde{v^{(j)}}(\pi_2) + H = \text{SW}[\omega]$ (using Eq. (6.9)) is worse than the choice of ω . For the allocation of $(S, \bar{\pi}_2)$, 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 6.6. *When $\bar{\pi}_1 = \emptyset$ or $\bar{\pi}_2 = \emptyset$, Claim 6.5 only requires the following small changes:*

- a. *When $\bar{\pi}_1 = \emptyset$ (i.e., $\pi_1 = [m]$), at any time $(T, \bar{\pi}_1)$ is a possible allocation declared in Claim 6.5, (T, R) for $R \subseteq \bar{T}$ is now also possible.⁵*
- b. *When $\bar{\pi}_2 = \emptyset$ (i.e., $\pi_2 = [m]$), at any time $(T, \bar{\pi}_2)$ is a possible allocation declared in Claim 6.5, (T, R) for $R \subseteq \bar{T}$ is now also possible.⁶*

Proof.

- a. This is because, due to the (strict) monotonicity of v we have $v(\pi_1) > v(\pi_2)$ and thus Eq. (6.8) tells us that $\Delta < 0$. Instead of choosing some sufficiently large H , we can choose $H = -\Delta$. It will make sure that $w(\emptyset) = w(\bar{\pi}_1) = 0$ while $w(\bar{\pi}_2) = -\Delta > 0$. The only place that we used H being sufficiently large, is where we declare that the only possible candidate allocation for $\text{VCG}(\cdot \sqcup w)$ is

⁵As a consequence, Claim 6.5a and Claim 6.5c still hold, but Claim 6.5b 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 $w \subseteq \bar{T}$.

⁶As a consequence, Claim 6.5b still holds, but Claim 6.5a and Claim 6.5c need small changes.

of the form $(S, \bar{\pi}_1)$ or $(S, \bar{\pi}_2)$. This is no longer true as we have to also consider (S, R) for $R \neq \bar{\pi}_1$ or $\bar{\pi}_2$. However, since $w(R) = 0$, $\text{SW}[(S, R)] = \text{SW}[(S, \emptyset)] = \text{SW}[(S, \bar{\pi}_1)]$. This means, allocation (S, R) will be possible *only if* $(S, \bar{\pi}_1)$ is possible.

- b. This is because, due to the weak monotonicity of $\widetilde{v^{(j^*)}}$ we have $\widetilde{v^{(j^*)}}(\pi_2) \geq \widetilde{v^{(j^*)}}(\pi_1)$ and thus Eq. (6.8) tells us that $\Delta > 0$. Instead of choosing some sufficiently large H , we can choose $H = 0$. It will make sure that $w(\emptyset) = w(\bar{\pi}_2) = 0$ while $w(\bar{\pi}_1) = \Delta > 0$. The only place that we used H being sufficiently large, is where we declare that the only possible candidate allocation for $\text{VCG}(\cdot \sqcup w)$ is of the form $(S, \bar{\pi}_1)$ or $(S, \bar{\pi}_2)$. This is no longer true as we have to also consider (S, R) for $R \neq \bar{\pi}_1$ or $\bar{\pi}_2$. However, since $w(R) = 0$, $\text{SW}[(S, R)] = \text{SW}[(S, \emptyset)] = \text{SW}[(S, \bar{\pi}_2)]$. This means, allocation (S, R) will be possible *only if* $(S, \bar{\pi}_2)$ is possible.

□

Now, we have some knowledge about what outcomes could be outputted by the VCG mechanism, on input $v \sqcup w$, and on bid $v^{(j)} \sqcup w$. We now come to the final part that is to show that Eq. (6.6) holds. We first compute the utilities in all three cases:

Claim 6.7. *If we choose $\theta(\pi_2) = K(\pi_2)^\top$ and $\theta(S) = K(S)^\perp$ for everything else (i.e., $S \neq \emptyset$ and $S \neq \pi_2$).*

- a. $U(\theta, \text{VCG}(v \sqcup w)) = K(\pi_2)^\top + H - \max_S w(S)$,
- b. $U(\theta, \text{VCG}(v^{(j^*)} \sqcup w)) \leq K(\pi_1)^\perp + H + \Delta - \max_S w(S)$ for every $j^* \in J$, and
- c. $U(\theta, \text{VCG}(v^{(j)} \sqcup w)) \leq K(\pi_2)^\top + H - \max_S w(S)$ for every $j \notin J$.

Proof.

- a. We have proved in Claim 6.5a that $(\pi_2, \bar{\pi}_2)$ is the only possible allocation in this case, and therefore $U(\theta, \text{VCG}(v \sqcup w)) = U(\theta, (\pi_2, \bar{\pi}_2)) = K(\pi_2)^\top + w(\bar{\pi}_2) - \max_S w(S) = K(\pi_2)^\top + H - \max_S w(S)$.
- b. We have proved in Claim 6.5b that $(T, \bar{\pi}_1)$ is the only possible allocation in this case, and therefore if $T \neq \pi_2$, we have $U(\theta, \text{VCG}(v^{(j^*)} \sqcup w)) = K(T)^\perp +$

$w(\overline{\pi_1}) - \max_S w(S) \leq K(\pi_1)^\perp + H + \Delta - \max_S w(S)$. (Here we used the weak monotonicity of K^\perp , i.e., $K(T)^\perp \leq K(\pi_1)^\perp$.)

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 Eq. (6.8), 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 $\text{SW}[(\pi_2, \overline{\pi_2})] = \text{SW}[(\pi_2, \overline{\pi_1})] - \Delta > \text{SW}[(\pi_2, \overline{\pi_1})]$. This indicates that $(\pi_2, \overline{\pi_1})$ will never be a possible outcome, giving a contradiction.

- c. We have proved in Claim 6.5c that $(T, \overline{\pi_2})$ is the only possible allocation in this case, and therefore $U(\theta, \text{VCG}(v^{(j^*)} \sqcup w)) \leq K(T)^\top + w(\overline{\pi_2}) - \max_S w(S) \leq K(\pi_2)^\top + w(\overline{\pi_2}) - \max_S w(S) = K(\pi_2)^\top + H - \max_S w(S)$. (Here we used the weak monotonicity of K^\top , i.e., $K(T)^\top \leq K(\pi_2)^\top$.)

We remark here that, in the case when $\overline{\pi_1} = \emptyset$ or $\overline{\pi_2} = \emptyset$, 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 6.8. *Eq. (6.6) is satisfied.*

Proof. We recall that Eq. (6.5) tells us that $v(\pi_2) - v(\pi_1) \leq K(\pi_2)^\top - K(\pi_1)^\perp$, but we have $v(\pi_2) - v(\pi_1) > \Delta$ in Eq. (6.8). This tells us that $K(\pi_2)^\top > K(\pi_1)^\perp + \Delta$.

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

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

while for every $j \notin J$,

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

The combination of them immediately implies Eq. (6.6) □

We recall that Eq. (6.6) gives a contradiction and says that v is an undominated strategy, and this ends the proof of the Bird Lemma 6.4, for Case 1.

Case 2

Suppose that the second inequality of Eq. (6.7) is broken, that is, $v(S') > \min_j \{v^{(j)}(S')\}$. Similar as Case 1, we let $J = \arg \min_j \{\widetilde{v}^{(j)}(S')\}$ be the set of minimizers, and let $j^* \in J$ be one of them. We can always choose some Δ such that

$$v(S') > \Delta > \widetilde{v}^{(j^*)}(S') , \quad (6.10)$$

and for every $j \notin J$:

$$\widetilde{v}^{(j)}(S') > \Delta . \quad (6.11)$$

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 $\overline{S'} \neq \emptyset$. We choose $H = 0$ if $\overline{S'} = \emptyset$.

Claim 6.9 (A variant of Claim 6.5). *If $\overline{S'} \neq \emptyset$,*

- a. *The allocation of VCG($v \sqcup w$) is $\omega = (S', \overline{S'})$*
- b. *For all $j^* \in J$, the allocation of VCG($v^{(j^*)} \sqcup w$) is $\omega = (\emptyset, [m])$.*
- c. *For all $j \notin J$, the allocation of VCG($v^{(j)} \sqcup w$) 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 bids w , if $T \notin \{\overline{S'}, [m]\}$, then $\text{SW}[(S, T)]$ does not contain the big term H and is thus smaller than any $\text{SW}[\omega]$ in all three cases. Therefore, we only need to consider outcomes of the form $(S, \overline{S'})$ and $(\emptyset, [m])$.

- a. In this case, $\text{SW}[\omega] = v(S') + H$. 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 $(\emptyset, [m])$ its social welfare $\text{SW}[(\emptyset, [m])] = \Delta + H < v(S') + H = \text{SW}[\omega]$, using Eq. (6.10). In sum, $\omega = (S', \overline{S'})$ must be the allocation of the VCG mechanism.
- b. In this case, $\text{SW}[\omega] = H + \Delta$. For the allocation of the form $(S, \overline{S'})$, $\text{SW}[(S, \overline{S}')] \leq \widetilde{v}^{(j^*)}(S) + H < H + \Delta = \text{SW}[\omega]$ (using Eq. (6.10)) is worse than the choice of

ω .

- c. In this case, $\text{SW}[\omega] = v^{(j)}(S') + H$. For the allocation of $(\emptyset, [m])$, we have that $\text{SW}[(\emptyset, [m])] = H + \Delta < v^{(j)}(S') + H = \text{SW}[\omega]$ (using Eq. (6.11)) 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 6.10 (A variant of Claim 6.6). *When $\overline{S'} = \emptyset$ (i.e., $S' = [m]$), Claim 6.9 only requires the following small changes:*

at any time $(T, \overline{S'})$ is a possible allocation declared in Claim 6.9, (T, R) for $R \subseteq \overline{T}$ is now also possible.⁷

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 $\text{VCG}(\cdot \sqcup w)$ is of the form $S, \overline{S'}$ or $(\emptyset, [m])$. This is no longer true as we have to also consider (S, R) for $R \neq \overline{S'}$ or $[m]$. However, since $w(R) = 0$, $\text{SW}[(S, R)] = \text{SW}[(S, \emptyset)] = \text{SW}[(S, \overline{S'})]$. 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 $v \sqcup w$, and on bid $v^{(j)} \sqcup w$. We now come to the final part that is to show that Eq. (6.6) holds. We first compute the utilities in all three cases:

Claim 6.11 (A variant of Claim 6.7). *If we choose $\theta(S) = K(S)^\top$ for everything non-empty S :*

- a. $U(\theta, \text{VCG}(v \sqcup w)) = K(S')^\top + H - \max_S w(S)$,
- b. $U(\theta, \text{VCG}(v^{(j^*)} \sqcup w)) = H + \Delta - \max_S w(S)$ for every $j^* \in J$, and
- c. $U(\theta, \text{VCG}(v^{(j)} \sqcup w)) \leq K(S')^\top + H - \max_S w(S)$ for every $j \notin J$.

Proof.

⁷As a consequence, Claim 6.9b still holds, but Claim 6.9a and Claim 6.9c need small changes.

- a. We have proved in Claim 6.9a that $(S', \overline{S'})$ is the only possible allocation in this case, and therefore $U(\theta, \text{VCG}(v \sqcup w)) = U(\theta, (S', \overline{S'})) = K(S')^\top + w(\overline{S'}) - \max_S w(S) = K(S')^\top + H - \max_S w(S)$.
(In the case when $\overline{S'} = \emptyset$, 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 6.9b that $(\emptyset, [m])$ is the only possible allocation in this case, and therefore $U(\theta, \text{VCG}(v^{(j^*)} \sqcup w)) = 0 + w([m]) - \max_S w(S) = H + \Delta - \max_S w(S)$.
- c. We have proved in Claim 6.9c that $(T, \overline{S'})$ is the only possible allocation in this case, and therefore $U(\theta, \text{VCG}(v^{(j)} \sqcup w)) \leq K(T)^\top + w(\overline{S'}) - \max_S w(S) \leq K(S')^\top + w(\overline{S'}) - \max_S w(S) = K(S')^\top + H - \max_S w(S)$.
(Here we used the weak monotonicity of K^\top , i.e., $K(T)^\top \leq K(S')^\top$. In the case when $\overline{S'} = \emptyset$, 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 6.12. *Eq. (6.6) is satisfied.*

Proof. We recall that Eq. (6.3) and Eq. (6.10) tell us that $\Delta < v(S') \leq K(S')^\top$. Now, for every $j^* \in J$,

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

while for every $j \notin J$,

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

The combination of them immediately implies Eq. (6.6)

□

We recall that Eq. (6.6) gives a contradiction and says that v is an undominated strategy, and this ends the proof of the Bird Lemma 6.4, for Case 2.

Case 3

Suppose that the second inequality of Eq. (6.7) is broken, that is, $v(S'') < \max_j \{v^{(j)}(S'')\}$. Similar as Case 1/2, we let $J = \arg \max_j \{\widetilde{v}^{(j)}(S'')\}$ be the set of maximizers, and let $j^* \in J$ be one of them. We can always choose some Δ such that

$$v(S'') < \Delta < \widetilde{v}^{(j^*)}(S'') , \quad (6.12)$$

and for every $j \notin J$:

$$\widetilde{v}^{(j)}(S'') < \Delta . \quad (6.13)$$

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 $\overline{S''} \neq \emptyset$. We choose $H = 0$ if $\overline{S''} = \emptyset$.

Claim 6.13 (A variant of Claim 6.5). *If $\overline{S''} \neq \emptyset$,*

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

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

- a. In this case, $\text{SW}[\omega] = H + \Delta$. 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 $\text{SW}[(S'', \overline{S''})] = v(S'') + H < H + \Delta = \text{SW}[\omega]$, using Eq. (6.12). In sum, $(\emptyset, [m])$ must be the allocation of the VCG mechanism.
- b. In this case, $\text{SW}[\omega] = \widetilde{v}^{(j^*)}(S'') + H$. For the allocation of $(\emptyset, [m])$, we have that $\text{SW}[(\emptyset, [m])] = H + \Delta < \widetilde{v}^{(j^*)}(S'') + H = \text{SW}[\omega]$ (using Eq. (6.12)) 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, $\text{SW}[\omega] = H + \Delta$. For the allocation of the form $(S, \overline{S''})$, $\text{SW}[(S, \overline{S''})] \leq \widetilde{v^{(j)}}(S) + H < H + \Delta = \text{SW}[\omega]$ (using Eq. (6.13)) is worse than the choice of ω .

□

Claim 6.14 (A variant of Claim 6.6). *When $\overline{S''} = \emptyset$ (i.e., $S'' = [m]$), Claim 6.13 only requires the following small changes:*

at any time $(T, \overline{S''})$ is a possible allocation declared in Claim 6.13, (T, R) for $R \subseteq \overline{T}$ is now also possible.⁸

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 $\text{VCG}(\cdot \sqcup w)$ is of the form $(S, \overline{S''})$ or $(\emptyset, [m])$. This is no longer true as we have to also consider (S, R) for $R \neq \overline{S''}$ or $[m]$. However, since $w(R) = 0$, $\text{SW}[(S, R)] = \text{SW}[(S, \emptyset)] = \text{SW}[(S, \overline{S''})]$. 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 $v \sqcup w$, and on bid $v^{(j)} \sqcup w$. We now come to the final part that is to show that Eq. (6.6) holds. We first compute the utilities in all three cases:

Claim 6.15 (A variant of Claim 6.7). *If we choose $\theta(S) = K(S)^\top$ for all non-empty S :*

- a. $U(\theta, \text{VCG}(v \sqcup w)) = H + \Delta - \max_S w(S)$,
- b. $U(\theta, \text{VCG}(v^{(j^*)} \sqcup w)) \leq H + K(S'')^\perp - \max_S w(S)$ for every $j^* \in J$, and
- c. $U(\theta, \text{VCG}(v^{(j)} \sqcup w)) = \Delta + H - \max_S w(S)$ for every $j \notin J$.

Proof.

- a. We have proved in Claim 6.13a that $(\emptyset, [m])$ is the only possible allocation in this case, and therefore $U(\theta, \text{VCG}(v \sqcup w)) = U(\theta, (\emptyset, [m])) = 0 + w(\overline{S''}) - \max_S w(S) = H + \Delta - \max_S w(S)$.

⁸As a consequence, Claim 6.13a and Claim 6.13c still hold, but Claim 6.13b needs small changes.

b. We have proved in Claim 6.13b that $(T, \overline{S''})$ is the only possible allocation in this case, and therefore $U(\theta, \text{VCG}(v^{(j^*)} \sqcup w)) \leq K(T)^\perp + w(\overline{S''}) - \max_S w(S) \leq K(S'')^\perp + w(\overline{S''}) - \max_S w(S) = K(S'')^\perp + H - \max_S w(S)$.

(Here we used the weak monotonicity of K^\perp , i.e., $K(T)^\perp \leq K(S'')^\perp$. In the case when $\overline{S''} = \emptyset$, 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 6.13c that $(\emptyset, [m])$ is the only possible allocation in this case, and therefore $U(\theta, \text{VCG}(v^{(j)} \sqcup w)) = 0 + w([m]) - \max_S w(S) = H + \Delta - \max_S w(S)$.

□

Corollary 6.16. *Eq. (6.6) is satisfied.*

Proof. We recall that Eq. (6.4) and Eq. (6.12) tell us that $\Delta > v(S'') \geq K(S'')^\perp$. Now, for every $j^* \in J$,

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

while for every $j \notin J$,

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

The combination of them immediately implies Eq. (6.6)

□

We recall that Eq. (6.6) gives a contradiction and says that v is an undominated strategy, and this ends the proof of the Bird Lemma 6.4, for Case 3.

6.3.4 The Hard Instance

In this subsection we will provide an example that is a two-player auction with $m \geq 2$ goods, and show that the social welfare may be broken exponentially, even though both players only consider undominated bidding strategies. In particular, we will provide approximate valuations K_1 and K_2 for the two players, undominated strategies $v_1 \in \text{UDed}(K_1)$ and $v_2 \in \text{UDed}(K_2)$, and at last show that in the game play of

$v_1 \sqcup v_2$, the allocation $\text{VCG}(v_1 \sqcup v_2)$ may be exponentially bad. Here we will use the Bird Lemma 6.4 twice to claim $v_1 \in \text{UDed}(K_1)$ and $v_2 \in \text{UDed}(K_2)$.

Before we start, let us consider a specific permutation π for over all $2^m - 1$ non-empty subsets of $[m]$, satisfying:

1. if $i < j$, then $\pi_i \not\subseteq \pi_j$;
2. $\pi_i = \overline{\pi_{2^m-1-i}}$; and
3. $\pi_{2^m-1} = [m]$.

For instance, when $m = 3$ we can let:

$$\pi = (\{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}) .$$

We choose some arbitrary positive constant x to start with:

Claim 6.17. $v_1 \in \text{UDed}(K_1)$ if we choose:

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

Proof. We only need to check that all assumptions in Bird Lemma 6.4 hold. Indeed, K_1^\perp and K_1^\top are 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 $v_1(S') = x \leq x(1 + \delta) = K_1(S')^\top$ and $v_1(S'') = x \geq x(1 - \delta) = K_1(S'')^\perp$. At last, we verify Eq. (6.5). If we just choose this specific π , we have:

$$\forall 1 \leq i \leq 2^m - 2, \quad v_1(\pi_i) - v_1(\pi_{i+1}) = -2\delta x = K_1(\pi_i)^\perp - K_1(\pi_{i+1})^\top ,$$

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

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

□

Now we choose some arbitrarily small constant $\varepsilon > 0$ and come to the second player:

Claim 6.18. $v_2 \in \text{UDed}(K_2)$ if we choose:

- K_2 to be such that

$$K_2(\pi_i) = \left[\left(\frac{(1+\delta)^i}{(1-\delta)^i} x - x - \frac{(1+\delta)^{i-1}}{(1-\delta)^i} \varepsilon \right) (1-\delta), \left(\frac{(1+\delta)^i}{(1-\delta)^i} x - x - \frac{(1+\delta)^{i-1}}{(1-\delta)^i} \varepsilon \right) (1+\delta) \right]$$

for all $i \in \{1, 2, \dots, 2^m - 2\}$, and $K_2(\pi_{2^m-1})$ being a set of a single value

$$K_2(\pi_{2^m-1}) = \{K_2(\pi_{2^m-2})^\top + x\} .$$

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

Proof. First, it is obvious that when ε is sufficiently small, $K_2(\pi_i)^\perp$ and $K_2(\pi_i)^\top$ are both positive. We now need to check that all assumptions in Bird Lemma 6.4 hold. Indeed, K_2^\perp, K_2^\top and v_2 are all strictly monotonic:

- $K_2^\perp(\pi_i) < K_2^\perp(\pi_j)$ for $i < j$,
- $K_2^\top(\pi_i) < K_2^\top(\pi_j)$ for $i < j$, and
- $v_1(\pi_i) < v_1(\pi_j)$ if $i < j$.

If we choose $S' = S'' = \pi_1$, we have $v_2(S') = 2\delta x - \varepsilon < (2\delta x - \varepsilon) \frac{1+\delta}{1-\delta} = K_2(S')^\top$ and $v_2(S'') = 2\delta x - \varepsilon = K_2(S'')^\perp$. At last, we verify Eq. (6.5). This time we choose the reverse permutation, that is, letting $\pi'_i = \pi_{2^m-i}$, and we have:

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

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

- for $i = 1$ and let $j = 2^m - 1 - i = 2^m - 2$:

$$\begin{aligned} v_2(\pi'_i) - v_2(\pi'_{i+1}) &= v_2(\pi_{j+1}) - v_2(\pi_j) = x \\ &= K_2(\pi_{j+1})^\perp - K_2(\pi_j)^\top = K_2(\pi'_i)^\perp - K_2(\pi'_{i+1})^\top , \text{ and} \end{aligned}$$

- for $i = 2^m - 1$, instead of computing the formula directly, we do something cleverer. If one sums up all inequalities in Eq. (6.5) for $1 \leq i \leq 2^m - 1$, he will

get

$$0 \geq \sum_j K_2(\pi'_j)^{\perp} - K_2(\pi'_j)^{\top}$$

which is obviously always true. However, we have already shown that when $1 \leq i \leq 2^m - 2$, the inequality is not only true but also tight, so the last one of them (which corresponds to $i = 2^m - 1$) has to be true.

□

6.3.5 Main Proof

Let the first two players have approximate valuations K_1 and K_2 , and bid v_1 and v_2 according to Claim 6.17 and Claim 6.18, and the rest of the players have zero approximate valuations and bid zero. We make the following observations.

- When players bid $v_1 \sqcup v_2 \sqcup 0 \sqcup \dots \sqcup 0$, the VCG mechanism will always pick allocation $\omega = ([m], \emptyset, \dots, \emptyset)$.

This is because, the social welfare computed using reported bids for ω is

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

while any other combination, if giving $\pi_i \neq \emptyset$ to player 1 and $\pi_{2^m-1-i} = \bar{\pi}_i$ to player 2, will result in a reported social welfare

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

smaller than ω . At last, if giving \emptyset to player 1 and $[m]$ to player 2, the reported social welfare is

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

also smaller than ω .

- Assume that we pick the true valuation $\theta_1 \in K_1$ such that $\theta_1(S) = x$ for all non-empty S , and $\theta_2 \in K_2$ such that $\theta_2(S) = K_2(S)^{\top}$. Of course, $TV_i(S) = 0$ for all $i > 2$.
- Now, the true social welfare on allocation ω is $SW(\theta, \omega) = \theta_1([m]) = x$.

- The maximum social welfare

$$\begin{aligned}
\text{MSW}(\theta) &\geq \theta_2([m]) = K_2(\pi_{2^{m-2}})^\top + x \\
&= \left(\frac{(1+\delta)^{2^m-2}}{(1-\delta)^{2^m-2}} x - x - \frac{(1+\delta)^{2^m-3}}{(1-\delta)^{2^m-2}} \varepsilon \right) (1+\delta) + x \\
&= \frac{(1+\delta)^{2^m-1}}{(1-\delta)^{2^m-2}} \left(x - \frac{\varepsilon}{1+\delta} \right) - \delta x .
\end{aligned}$$

- In sum, we have shown that the fraction of the social welfare guaranteed in this case is

$$\frac{\text{SW}(\theta, \omega)}{\text{MSW}(\theta)} = \frac{1}{\frac{(1+\delta)^{2^m-1}}{(1-\delta)^{2^m-2}} \left(1 - \frac{\varepsilon}{x+\delta x} \right) - \delta} ,$$

and since we can choose $\varepsilon > 0$ arbitrarily small, the social welfare guarantee of the VCG mechanism is at most

$$\frac{1}{\frac{(1+\delta)^{2^m-1}}{(1-\delta)^{2^m-2}} - \delta} .$$

At last, we slightly weaken the bound that we have just proved, to provide a cleaner fraction:

$$\frac{1}{\frac{(1+\delta)^{2^m-1}}{(1-\delta)^{2^m-2}} - \delta} = \frac{(1-\delta)^{2^m-2}}{(1+\delta)^{2^m-1} - \delta(1-\delta)^{2^m-2}} \leq \frac{(1-\delta)^{2^m-2}}{(1+\delta)^{2^m-2}} = \left(\frac{1-\delta}{1+\delta} \right)^{2^m-2} .$$

■

Chapter 7

General Games

In this chapter we formally generalize our notions in the single-good or multiple-good auction cases, into any general game where players have approximate *types*.

7.1 Game: Context and Mechanism

The most basic notion is that of a pre-context: it keeps track of the number of players, the set of possible types for each of the players, the set of all possible outcomes, the utility function for each of the players, and a promise on each player’s “quality of approximation” about his own type. Formally:

Definition 7.1. A pre-context is a tuple $\text{PreC} = ([n], \Theta, \Omega, U, Q)$ such that:

- $[n] = \{1, 2, \dots, n\}$ is a set of players;
- $\Theta = (\Theta_i)_{i \in [n]}$ where each Θ_i is a set of types for player i ;
- Ω is a set of outcomes;
- $U = (U_i)_{i \in [n]}$ where each $U_i: \Theta_i \times \Omega \rightarrow \mathbb{R}$ is the utility function of player i ; and
- $Q = (Q_i)_{i \in [n]}$ where each Q_i , a subset of 2^{Θ_i} , indicates the approximation quality of player i .

Remark 7.2. A “new” component in Definition 7.1 (as compared to the usual definition of a pre-context) is the *approximation quality* Q that players are promised to have: for each player i , Q_i is an explicit collection of subsets of Θ_i ; each subset in

Q_i is a possible *approximate type* for player i . Our definition reduces to the usual definition of a pre-context whenever each Q_i is the collection of all singletons in 2^{Θ_i} (i.e., the only “allowed” approximation is the perfect approximation).

A pre-context becomes a context when it is augmented with a *specific* set of players, i.e., a specific profile of approximate types and true (but secret) types:

Definition 7.3. *A context is a tuple $C = ([n], \Theta, \Omega, U, Q, K, \theta)$ such that:*

- $\text{PreC} = ([n], \Theta, \Omega, U, Q)$ is a pre-context;
- $K = (K_i)_{i \in [n]}$ where each K_i , a set in Q_i , is the approximate type of player i ;
and
- $\theta \in K$ where each θ_i is the true type of player i .

Notation. We emphasize that by augmenting appropriate $K \in Q$, and appropriate $\theta \in K$, a pre-context PreC actually specifies a *class* of contexts (just like our $\mathcal{C}_{n,B}^\delta$ in the single-good auction case). We therefore in this chapter allows quantifiers like “ $\forall C \in \text{PreC}$ ”.

As usual, a mechanism specifies the players’ strategies and how these strategies determine outcomes:

Definition 7.4. *A mechanism for players $[n]$ and outcome set Ω is a tuple $M = (S, F)$ such that:*

- $S = (S_i)_{i \in [n]}$ where each S_i is the set of strategies for player i ; and
- $F: S \rightarrow \Omega$ is the (possibly probabilistic) outcome function.

We say that M is for a pre-context PreC if $[n]$ and Ω respectively match the set of players and set of outcomes in PreC .

As usual, we will denote strategies with Latin letters (such as s and t), and mixtures of strategies with Greek letters (such as σ and τ).

Finally, we state the knowledge model considered in this thesis:

Definition 7.5. *In the incomplete information knowledge model, when considering a context $C = ([n], \Theta, \Omega, U, Q, K, \theta)$ and a mechanism M ,*

- the pre-context $([n], \Theta, \Omega, U, Q)$ is common knowledge;
- each K_i is known to player i (and no other player);
- each θ_i is secret to player i (and all other players); and
- the mechanism M is common knowledge.

7.2 Dominance

A natural comparison to make between two actions is establishing some kind of *dominance* relation. We discuss the three natural definitions of dominance, whose intuitive meaning we now summarize:

- “never worse than” — an action σ_i of player i *very-weakly dominates* another action σ'_i if σ_i is never worse than σ'_i , when considering all other players’ possible actions;
- “never worse than, and at least once better” — an action σ_i of player i *weakly dominates* another action σ'_i if σ_i is never worse than σ'_i , when considering all other players’ possible actions, and σ_i is better than σ'_i for at least one choice of other players’ actions; and
- “always better” — an action σ_i of player i *strictly dominates* another action σ'_i if σ_i is always better than σ'_i , when considering all other players’ possible actions.

Formally:

Definition 7.6. *Given mechanism $M = (S, F)$ for pre-context $\text{PreC} = ([n], \Theta, \Omega, U, Q)$, and a player $i \in [n]$ with approximate type $K_i \in Q_i$, for his two (possibly mixed) strategies $\sigma_i \in \Delta(S_i)$ and $\sigma'_i \in \Delta(S_i)$, we say that:*

- σ_i **very-weakly dominates** σ'_i , in symbols $\sigma_i \succeq_{i, K_i} \sigma'_i$, if:

$$\forall \theta_i \in K_i, \forall \tau_{-i} \in S_{-i}, \mathbb{E}U_i(\theta_i, F(\sigma_i \sqcup \tau_{-i})) \geq \mathbb{E}U_i(\theta_i, F(\sigma'_i \sqcup \tau_{-i})) .$$

- σ_i weakly dominates σ'_i , in symbols $\sigma_i \overset{w}{\succ}_{i, K_i} \sigma'_i$, if:

$$\forall \theta_i \in K_i, \forall \tau_{-i} \in S_{-i}, \mathbb{E}U_i(\theta_i, F(\sigma_i \sqcup \tau_{-i})) \geq \mathbb{E}U_i(\theta_i, F(\sigma'_i \sqcup \tau_{-i})) , \text{ and} \quad (7.1)$$

$$\exists \theta_i \in K_i, \exists \tau_{-i} \in S_{-i}, \mathbb{E}U_i(\theta_i, F(\sigma_i \sqcup \tau_{-i})) > \mathbb{E}U_i(\theta_i, F(\sigma'_i \sqcup \tau_{-i})) . \quad (7.2)$$

- σ_i strictly dominates σ'_i , in symbols $\sigma_i \overset{s}{\succ}_{i, K_i} \sigma'_i$, if:

$$\forall \theta_i \in K_i, \forall \tau_{-i} \in S_{-i}, \mathbb{E}U_i(\theta_i, F(\sigma_i \sqcup \tau_{-i})) > \mathbb{E}U_i(\theta_i, F(\sigma'_i \sqcup \tau_{-i})) .$$

(It is a simply exercise to see that those definitions remain the same if one replaces “ $\tau_{-i} \in S_{-i}$ ” by “ $\tau_{-i} \in \Delta(S_{-i})$ ”, incorporating mixed strategies for the rest of the players.)

Remark 7.7. We have defined the approximate type K_i of a player to be an *arbitrary* subset of Θ_i , because we certainly do not want to place any restrictions on the structure of a player’s lack of knowledge about his own true type (e.g., by assuming that it is an interval).

In the setting of a (known) single-parameter domain, however, where $\Theta_i = [0, B]$ for some large B , a player’s reasoning can be greatly simplified. Specifically, a player with approximate type K_i could equivalently reason in the world $K'_i = [\min K_i, \max K_i]$ or in the world $K''_i = \{\min K_i, \max K_i\}$. This is an easy consequence of Definition 7.6: there is only one parameter to play with on two sides of the inequalities, so we only need to care about extreme points.

For instance, in a single good auction, a player with an approximate valuation $K_i = \{90, 91, 92, 100\}$ might as well imagine he is in the world $K'_i = [90, 100]$ or in the world $K''_i = [90, 100]$. In other words, only the extremal points of K_i matter for the purpose of determining dominance relations, and all other points are irrelevant.

7.3 Solution Concepts

The three notions of dominance from Definition 7.6 should give rise to three corresponding sets of “undominated” strategies and three corresponding sets of “domi-

nant” strategies, both in the approximate player type case. However, following the convention of game theory, we here only define and study the set of very-weakly-dominant strategies, and the set of not-weakly-dominated strategies.

Definition 7.8. *Given mechanism $M = (S, F)$ for pre-context $\text{PreC} = ([n], \Theta, \Omega, U, Q)$, and a player $i \in [n]$ with approximate type $K_i \in Q_i$, then*

- *The set of **very-weakly-dominant** (or simply **dominant**) strategies of i with respect to K_i is defined as*

$$\text{Dnt}_i(K_i) \stackrel{\text{def}}{=} \left\{ s_i \in S_i : \forall \tau_i \in \Delta(S_i), s_i \succeq_{i, K_i} \tau_i \right\} .$$

- *The set of **not-weakly-dominated** (or simply **undominated**) strategies of i with respect to K_i is defined as*

$$\text{UDed}_i(K_i) \stackrel{\text{def}}{=} \left\{ s_i \in S_i : \nexists \tau_i \in \Delta(S_i) \text{ s.t. } \tau_i \succ_{i, K_i}^w s_i \right\} .$$

Note that all the above definitions reduce to the usual ones in the case where all the players’ types are exact (i.e., K_i is a singleton).

Notation. Whenever we ignore the subscript we denote the Cartesian product of the set among players. For instance, $\text{UDed}(K) \stackrel{\text{def}}{=} \text{UDed}_1(K_1) \times \cdots \times \text{UDed}_n(K_n)$.

Next, we prove that whenever there is at least one very-weakly-dominant strategy, then every strategy that is undominated is also a very-weakly-dominant strategy. We will invoke this lemma at times for the purpose of understanding the meaningfulness of certain solution concepts. See Remark 7.12.

Lemma 7.9. *Fix a game \mathcal{G} . For every player $i \in [n]$,*

$$\text{Dnt}_i(K_i) \neq \emptyset \longrightarrow \text{Dnt}_i(K_i) = \text{UDed}_i(K_i) .$$

Proof. The set inclusion $\text{Dnt}_i(K_i) \subseteq \text{UDed}_i(K_i)$ easily follows from the definition of the two sets. The reverse set inclusion, $\text{Dnt}_i(K_i) \supseteq \text{UDed}_i(K_i)$, is the interesting one. So consider a strategy $s_i \in \text{UDed}_i(K_i)$. Suppose by way of contradiction that there exists a strategy $\sigma_i \in \Delta(S_i)$ for which it is *not* the case that $s_i \succeq_{i, K_i} \sigma_i$, i.e.,

$$\exists \theta_i \in K_i, \exists \tau_{-i} \in S_{-i} : \mathbb{E}U_i(\theta_i, F(s_i \sqcup \tau_{-i})) < \mathbb{E}U_i(\theta_i, F(\sigma_i \sqcup \tau_{-i})) . \quad (7.3)$$

Let σ_i^* be any strategy in $\text{Dnt}_i(K_i)$, so that, in particular, we know that $\sigma_i^* \succeq_{i,K_i} \sigma_i$, and thus

$$\mathbb{E}U_i(\theta_i, F(\sigma_i \sqcup \tau_{-i})) \leq \mathbb{E}U_i(\theta_i, F(\sigma_i^* \sqcup \tau_{-i})) . \quad (7.4)$$

We also know that $\sigma_i^* \succeq_{i,K_i} s_i$, which, combined with Eq. (7.3) and Eq. (7.4), yields that $\sigma_i^* \succ_{i,K_i}^w s_i$, contradicting the fact that $s_i \in \text{UDed}_i(K_i)$. We conclude that $s_i \in \text{Dnt}_i(K_i)$, as desired. \square

7.4 Implementation

We can now state what it means for a mechanism to “implement” a social property relative to a pre-context (i.e., a certain class of contexts) in the approximate-type world. A social property is simply a predicate over the players’ types and an outcome distribution, indicating which possible outcome is desirable:

Definition 7.10. *A social property is a function $\Pi: \Theta \times \Delta(\Omega) \rightarrow \{0, 1\}$.*

Definition 7.11 (Implementation). *Let M be a mechanism for $\text{PreC} = ([n], \Theta, \Omega, U, Q)$ and Π a social property. We say that:*

- M implements Π in dominant strategies (with respect to PreC) if

$$\forall K \in Q, \exists s \in \text{Dnt}(K), \forall \theta \in K : \Pi(\theta, F(s)) = 1 .$$

- M implements Π in undominated strategies (with respect to PreC) if

$$\forall K \in Q, \forall s \in \text{UDed}(K), \forall \theta \in K : \Pi(\theta, F(s)) = 1 .$$

Remark 7.12.

- We actually choose a *worst-case perspective*. That is, we require that the property Π hold at *every possible type* $\theta \in K$.
- Following the convention, we put an existential quantifier to s for dominant-strategy implementation, but a universal quantifier for undominated-strategy implementation. Although being weaker in the former case, to be proved in

Section 7.5, it is actually equivalent to dominant-strategy-truthful mechanisms in the approximate-type world by revelation principle.

- If one also put a universal quantifier on s for dominant-strategy implementation, this safer notion, according to Lemma 7.9 is equivalent to implementation in undominated strategies. (But, of course, the reverse does not always hold.)

7.5 New Revelation Principle

We prove an analogue of the Revelation Principle in the setting of approximate player types. A version of the classical Revelation Principle states that, when considering partial implementations in dominant strategies, it suffices to consider dominant-strategy-truthful mechanisms where strategy sets are identical to type spaces Θ , and reporting the true type is a (very-weakly-)dominant strategy.

In our setting of approximate types, we will construct a mechanism where the strategies are identical to approximate types Q .¹ We will invoke the Revelation Principle in our impossibility results for implementation in very-weakly-dominant strategies.

Lemma 7.13 (Revelation Principle). *Given a mechanism $M = (S, F)$ for $\text{PreC} = ([n], \Theta, \Omega, U, Q)$, and social choice function $W : Q \rightarrow \Delta(\Omega)$ such that,*

$$\forall K \in Q, \exists s \in \text{Dnt}(K) \text{ s.t. } F(s) = W(K) ,$$

then there exists a direct mechanism $M' = (S', F')$ for which $S' = Q$ and

$$\forall K \in Q, K \in \text{Dnt}'(K) \text{ and } F'(K) = W(K) .$$

In other words, the direct mechanism M' is a (very-weakly-)dominant-strategy-truthful one that guarantees reporting the truth K is (very-weakly-)dominant, and yields the same outcome. Here Dnt' is the set of dominant strategies in M' .

¹Here is one place where it becomes important for pre-contexts to contain the component Q promising what kinds of approximate types the players will have. If we had not introduced the additional component Q , and defined the set of strategies in the revealed mechanism to be 2^{Θ^i} for player i , the Revelation Principle would no longer hold.

Proof. There exists a function $f: Q \rightarrow \Delta(S)$ such that $f(K)$ is the lexicographically first strategy $s \in \text{Dnt}(K)$ for which $M(s) = W(K)$. Consider the mechanism $M' = (S', F')$ for which $S' \stackrel{\text{def}}{=} Q$ and $F'(K) \stackrel{\text{def}}{=} F(f(K))$ for every $K \in Q$.

Fix any $K \in Q$ and suppose that $\exists \sigma \in \text{Dnt}(K)$ for which $M(\sigma) = W(K)$, and let it be the lexicographically first one. By construction, $F'(K) = F(f(K)) = F(\sigma) = W(K)$. We are left to prove $K \in \text{Dnt}'(K)$. Indeed, consider any alternative strategy $K'_i \in Q$ for player i . Then, for every $\theta_i \in K_i$ and every $K_{-i} \in Q_{-i}$,

$$\begin{aligned} \mathbb{E}U(\theta_i, F'(K_i \sqcup K_{-i})) &= \mathbb{E}U(\theta_i, F(f(K_i \sqcup K_{-i}))) \\ &\geq \mathbb{E}U(\theta_i, F(f(K'_i \sqcup K_{-i}))) = \mathbb{E}U(\theta_i, F'(K'_i \sqcup K_{-i})) , \end{aligned}$$

where the inequality comes from the fact that $f(K_i \sqcup K_{-i}) \in \text{Dnt}(K)$. This completes the proof. \square

Corollary 7.14. *Fix any mechanism $M = (S, F)$ for $\text{PreC} = ([n], \Theta, \Omega, U, Q)$ that partially implements some social property $\Pi: \Theta \times \Delta(\Omega) \rightarrow \{0, 1\}$ in dominant strategies. We can construct a direct dominant-strategy-truthful (DST) mechanism M' such that reporting the true K is dominant and implements Π .*

Indeed, an analogue of Lemma 7.13 also holds for implementation in ex-post Nash, and therefore all of our impossibility results in Theorem 2b and Theorem 3b can be extended to implementation in ex-post Nash equilibria. (See, for example, Corollary 9.26 in [26] for the idea.)

Chapter 8

Universality of Approximate Types

While accounting for the uncertainty that players have about their own types by introducing a set-theoretic approximate type (as we did throughout this thesis) is very well motivated, in this chapter we argue that for *any* game, *any* reasonable (including Bayesian or Knightian) model for player type uncertainty is *equivalent* to our set-theoretic approximate-type model, at least for the setting when the type space Θ is “convex” (which is the case in almost all natural problem domains; see Remark 8.1 below).

More concretely, we describe more and more general models of player type uncertainty, where the most general one is an infinite “hierarchy type uncertainty” (modeling uncertainty, uncertainty about uncertainty, and so on) consisting of alternating levels of sets and distributions. Our set-theoretic approximate type is merely one of the two types of models in the first level of this hierarchy. For convex Θ , we show how this hierarchy “collapses” to our model, thus proving that it is *universal* as far as modeling type uncertainty of players is concerned. Details follow.

No Type Uncertainty. Let us first recall what it means for a player to know his own type *exactly*: if player i knows that his type is θ_i (and is completely sure about it!), when he sees an outcome $\omega \in \Omega$, his utility is by definition the single real $U_i(\theta_i, \omega) \in \mathbb{R}$.

Approximate Type Uncertainty. One way of relaxing the No Type Uncertainty

model into a safer model is to consider a set-theoretic approximate type model, which is the model that we introduced in Section 7.1 (and whose universality we are going to prove). Recall, from Definition 7.3, that a context captures the internal knowledge of player i via an *approximate type* K_i , which is a *set* that is promised to contain the true type θ_i of player i . Therefore, when player i sees an outcome $\omega \in \Omega$, his utility is a *set* of reals $\{U_i(\theta_i, \omega)\}_{\theta_i \in K_i}$, and, because player i knows that $\theta_i \in K_i$, he also knows that $U_i(\theta_i, \omega)$ is in his “utility set”.

Bayesian Type Uncertainty. Another way of relaxing the No Type Uncertainty model into a safer model is to consider an (individual) Bayesian type model, where each player i does not only know a set K_i containing his true type θ_i but also knows a *distribution* $D_i^* \in \Delta(\Theta_i)$ over K_i assigning probabilities to the possible values for θ_i . Let us recall the following widely-accepted assumption about the behavior of players:

Assumption 1 (Bernoulli¹). A player seeks to maximize his own expected utility.

Under Bernoulli’s Assumption, when player i sees an outcome $\omega \in \Omega$, his utility is the single real value $\mathbb{E}_{\theta_i \sim D_i}[U_i(\theta_i, \omega)]$.

While this model is formally “safer” than the No Type Uncertainty model, it can be collapsed to it whenever Θ_i is “convex”; indeed, in such a case, we can write $\mathbb{E}_{\theta_i \sim D_i}[U_i(\theta_i, \omega)] = U_i(\theta'_i, \omega)$ with θ'_i defined informally as “ $\mathbb{E}_{\theta_i \sim D_i}[\theta_i]$ ”, so that player i can equivalently take θ'_i as his exact true type. (See Theorem 7 for the formal statement.)

Knightian Type Uncertainty. We could combine the previous two relaxations into an even safer model, an *Knightian* type model, where each player i knows a *set of distributions* \mathcal{D}_i that is promised to contain the true distribution D_i^* from which his true type θ_i is drawn. In other words, a player i may individually know that his own true type is drawn from some distribution D_i^* in some class of distributions \mathcal{D}_i , without being sure of *which* distribution in \mathcal{D}_i is the right one.

For example, while a player may not know the exact distribution from which his true type is drawn, he may still know all kinds of probabilistic information about his

¹A theory on expected utility was first introduced by Daniel Bernoulli ([3] in Latin; see [4] for an English translation), where he used the terminology of *moral expectation*.

own true type, and these estimates restrict the class of possible distributions from which his true type may be drawn. The better the probabilistic information of the player, the “smaller” the set \mathcal{D}_i .

This formalization of uncertainty was also studied by Bewley [5], as a rigorous expression to Knight’s distinction between risk and uncertainty [20]: a random variable is risky if a distribution is known, while uncertain if the distribution is unknown.

Again invoking Bernoulli’s Assumption, when player i sees an outcome $\omega \in \Omega$, his utility is a set $\{\mathbb{E}_{\theta_i \sim D_i}[U_i(\theta_i, \omega)]\}_{D_i \in \mathcal{D}_i}$, and, because player i knows that $D_i^* \in \mathcal{D}_i$, he also knows that $\mathbb{E}_{\theta_i \sim D_i^*}[U_i(\theta_i, \omega)]$ is in his “utility set”.

As before, while this model is formally “safer” than the Approximate Type Uncertainty model, it can be collapsed to it whenever Θ_i is “convex”; indeed, in such a case, we can write “ $\mathbb{E}_{\theta_i \sim D_i}[U_i(\theta_i, \omega)] = U_i(\mathbb{E}_{\theta_i \sim D_i}[\theta_i], \omega)$ ”, and thus $\{\mathbb{E}_{\theta_i \sim D_i}[U_i(\theta_i, \omega)]\}_{D_i \in \mathcal{D}_i} = \{U_i(\mathbb{E}_{\theta_i \sim D_i}[\theta_i], \omega)\}_{D_i \in \mathcal{D}_i}$, so that player i can equivalently take $K_i := \{\mathbb{E}_{\theta_i \sim D_i}[\theta_i]\}_{D_i \in \mathcal{D}_i}$ as his approximate type. (See Theorem 7 for the formal statement.)

Hierarchy Type Uncertainty. Why stop at Knightian uncertainty? Note the pattern:

- Approximate Type Uncertainty \equiv set of types;
- Bayesian Type Uncertainty \equiv distribution over types;
- Knightian Type Uncertainty \equiv set of distribution over types.

How about a distribution over sets of types? How about a set of distributions over sets of types? How about a distribution over distributions over sets of distributions over types?

In the most general case, there is a *type uncertainty hierarchy*; all the models we have considered so far are special cases of this hierarchy. More precisely, a player’s knowledge about his own true type θ_i can be recursively defined as a (rooted) tree in which each internal node is either a “set node” (indicating that the player type could be any one of its subtrees) or a “distribution node” (indicating an exact probabilistic distribution over its subtrees), and the leaves of the tree contain elements from Θ_i (and are thus the candidates for θ_i).

Without loss of generality, we can require the set nodes and distribution nodes to

be at alternating levels in the hierarchy. This is because, if there is a set of sets (i.e., a set node all of whose children are also set nodes) or a distribution over distributions (i.e., a distribution node all of whose children are also distribution nodes), the two levels of the tree collapse into one. A more sophisticated case is for instance a set node with two children: a set node and a distribution node. In this case, we can also absorb the set node subtree into the root.

We are unaware of any reasonable type uncertainty model (suitable for the setting of incomplete information) that cannot be defined within our type uncertainty hierarchy.

For instance, the Approximate Type Uncertainty model is a depth-1 tree, with the only internal node (the root) being a set node (see Figure 8-1a). The Bayesian Type Uncertainty model is also a depth-1 tree, but with a distribution node as the root (see Figure 8-1b). The Knightian Type Uncertainty model is a depth-2 tree, with its root being a set node and each middle-layer-child being a distribution node (see Figure 8-1c).

Equivalence. Now we are going to prove that any tree uncertainty model is equivalent to our set-theoretic approximate type model, at least for the setting that the type space is “convex”:

Theorem 7. *If the set of functions $\{U_i(\theta_i, \cdot): \Omega \rightarrow \mathbb{R}\}_{\theta_i \in \Theta_i}$ is convex, then any Hierarchy Type Uncertainty model is equivalent to the Approximate Type Uncertainty model.*

We deduce that it *only suffices to study the approximate type model*, which is simpler to work with anyways!

Proof. Without loss of generality, we assume that the tree is balanced, since we can always add set nodes with one child, and distribution nodes with one child. We have also argued that we can assume that the tree has alternating structure: the child of a set (resp., distribution) node is either a distribution (resp., set) node or a leaf.

As a warm-up, let us first recall the equivalence for the Bayesian Type Uncertainty model. For player i that knows $D_i \in \Delta(\Theta_i)$, by Bernoulli’s Assumption, he is going to





Figure 8-1: Three basic models of type uncertainty in our hierarchy.

maximize his expected utility $\mathbb{E}_{\theta_i \sim D_i}[U_i(\theta_i, \omega)]$. By the convexity assumption, there exists some θ'_i such that the function $U_i(\theta'_i, \cdot) \equiv \mathbb{E}_{\theta_i \sim D_i}[U_i(\theta_i, \cdot)]$, this means it is equivalent from player i 's perspective to have an exact type $\theta_i = \theta'_i$, no matter which outcome is chosen. We call this θ'_i the expected type.

In general, the above proof shows that if the lowest level of the tree is a distribution node, we can replace such node along with its children (leaves) by the expected type, reducing the depth of the entire tree by 1. As a consequence, for instance, the Bayesian Type Uncertainty model in Figure 8-1b collapses to the No Type Uncertainty model; while the Approximate Bayesian Type Uncertainty model in Figure 8-1c collapses to the Approximate Type Uncertainty model in Figure 8-1a.

Next, we are going to prove that whenever we have a distribution of sets of types, this collapses into one single set of types. In other words, a depth-2 sub-tree at the lowest level of the whole tree, if it is a distribution node with children being set nodes, can be replaced by a set node directly, again reducing the depth of the entire tree by

1.

Suppose that we have a distribution \mathcal{K}_i over $\{K_i^{(1)}, K_i^{(2)}, \dots\}$, while $K_i^{(j)}$ is the actual set with probability $p_i^{(j)}$. Player i knows that, with probability $p_i^{(j)}$, his utility is among the set $K_i^{(j)}$. By Bernoulli's Assumption, we can compute the expected utility of player i by taking an arbitrary element from each set $K_i^{(j)}$, and then compute the weighted average according to $p_{(j)}$. This is again a set:

$$\left\{ \sum_j p_i^{(j)} \theta_i^{(j)} \mid \theta_i^{(j)} \in K_i^{(j)} \right\} .$$

In sum, if we continue to collapse the bottom two levels of the tree, and combine two consecutive levels of sets, we will eventually get a depth-1 tree with a single set node.

In other words, we can replace the knowledge of player i with a set of types computed by collapsing his tree; these set of types can be modeled within the Approximate Type Uncertainty model, concluding the proof of the theorem. \blacksquare

Remark 8.1. In many problem domains, the convexity of $\{U_i(\theta_i, \cdot): \Omega \rightarrow \mathbb{R}\}_{\theta_i \in \Theta_i}$ is easily satisfied. For instance, in single-good auctions, θ_i is player i 's valuation to the good on sale, and

$$U_i(\theta_i, \omega) := \begin{cases} \theta_i - P, & \text{if player } i \text{ wins with price } P \\ 0, & \text{if some other player wins} \end{cases} .$$

The convexity condition in this case is equivalent to saying that if a player can have a valuation of both a and b , then his valuation may also be any real number in the interval $[a, b]$.

According to Theorem 7, if player i knows an exact Bayesian distribution D_i for his true type θ_i , he can simply imagine that $\theta'_i := \mathbb{E}_{\theta_i \sim D_i}[\theta_i]$ is his exact true type. Therefore if the auction is the second-price auction, bidding θ'_i is a dominant strategy.²

Of course, one can also find less natural utility functions that do not satisfy the convexity property. For instance Sandholm [30] also considers the Bayesian Type Uncertainty model for the single-good auction, but with the utility unconventionally

²In our language, this is a very-weakly dominant strategy.

defined as:

$$U_i(\theta_i, \omega) := \begin{cases} \theta_i - P, & \text{if player } i \text{ wins with price } P \leq \theta_i \\ 2(\theta_i - P), & \text{if player } i \text{ wins with price } P > \theta_i \\ 0, & \text{if some other player wins} \end{cases} .$$

In this case the convexity is no longer satisfied. And, indeed, for example, in a second-price auction, if player i has an exact distribution $D_i = U(0, 1)$ that is uniform between 0 and 1, bidding the expectation 0.5 is no longer a dominant strategy for him.

Chapter 9

Conclusion

We believe that assuming uncertainty about the players' internal knowledge is mandatory if we want to develop a realistic theory of strategic human interaction. Of course modeling uncertainty via a probability distribution is extremely effective, if we can lay hands on the true distribution. But often a player may not be able to write down the knowledge of his valuations in such a precise form. Modeling it in an approximate, and actually in a purely set theoretic way is certainly safer. But: is it useful? That is, can we leverage it in mechanism design?

The answer would obviously be no if there is no restriction on such sets, and this is why we consider δ -approximate valuations for some constant $\delta \in [0, 1)$, and study the performance guarantee as a function of δ . To the best of our knowledge, this is the first time such tradeoff between player's internal knowledge and mechanism's performance is studied in literature, and we have provided tight bounds for single-good auctions.


For combinatorial auctions, our impossibility results are very strong but they should not be interpreted as bad news for mechanism design with approximate types. Ultimately, they will guide us towards better approaches. For instance, by allowing players to have external knowledge about other players, one can still guarantee moderate social welfare or revenue benchmark, as observed by the same author [7]. Such results are actually very encouraging, knowing that truly combinatorial auctions are truly challenging.

Before this work, almost all set-theoretic (a.k.a. Knightian) analysis of player's


type uncertainty has been immersed in protected waters of decision theory, which studies purely from a player's perspective about how he will possibly behave when he is facing a set of candidate types (or distributions). However, we believe that it is time to navigate the more open ones of *mechanism design*, even though players have such *approximate types* and do not necessarily know which candidate is better.

Appendix A


Performance Diagrams




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



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



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



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

Figure A-1: Performance of our optimal mechanism

In Figure A-1 we compare the social welfare guarantees of:

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

In Figure A-1a and Figure A-1b we compare ε versus δ , and in Figure A-1c and Figure A-1d we compare ε versus n . The green data, our mechanism, is always better (at times significantly) than the other two mechanisms.

Appendix B

Our Optimal Mechanism $M_{\text{opt}}^{(\delta)}$

In this section we provide a concise description to our optimal mechanism in Theorem 3a for a single-good auction with n players, valuation bound B and approximation accuracy δ . We first construct the following allocation function:

Definition B.1. For every $\delta \in (0, 1)$, and let $D_\delta \stackrel{\text{def}}{=} \left(\frac{1+\delta}{1-\delta}\right)^2 - 1 > 0$. We define the function $f^{(\delta)}: [0, B]^{[n]} \rightarrow [0, 1]^{[n]}$ as follows:

- for every $z = (z_1, \dots, z_n) \in [0, B]^{[n]}$ such that $z_1 \geq z_2 \geq \dots \geq z_n$, let $n^* \in \{1, 2, \dots, n\}$ be the index in $[n]$ (which exists and is unique) such that

$$z_1 \geq \dots \geq z_{n^*} > \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta} \geq z_{n^*+1} \geq \dots \geq z_n ,$$

Then set

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

- for other z , define $f^{(\delta)}$ by extending it symmetrically.

The (code for the) outcome function of our mechanism $M_{\text{opt}}^{(\delta)}$ is:

Code for outcome function of $M_{\text{opt}}^{(\delta)}$

public parameter: $\delta \in (0, 1)$

inputs: $v_1, \dots, v_n \in \{0, 1, \dots, B\}$

output: (i, P) , where $i \in [n] \cup \{\perp\}$ is the winning player and $P \in \mathbb{R}^{[n]}$ is the price profile

pseudocode:

1. Draw r uniformly at random in $[0, 1]$.
2. (Define $f_0^{(\delta)} \stackrel{\text{def}}{=} 0$.)
3. If there exists $i \in [n]$ such that $\sum_{j=0}^{i-1} f_j^{(\delta)}(v) < r \leq \sum_{j=0}^i f_j^{(\delta)}(v)$:
 - Compute $P_i \stackrel{\text{def}}{=} v_i - \frac{\int_0^{v_i} f_i^{(\delta)}(z \perp v_{-i}) dz}{f_i^{(\delta)}(v)}$, and $P_j = 0$ for $j \neq i$, and output (i, P) .
4. Otherwise, output $(\perp, (0, \dots, 0))$. (No player is assigned the good.)

We note that our mechanism can be tweaked to make sure that the good is always assigned to some player. But the proof is more involved than it already is, and we leave it to a future version of this paper.

Bibliography

- [1] Robert J. Aumann. Utility theory without the completeness axiom. *Econometrica*, 30(3):445–462, July 1962.
- [2] 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.
- [3] Daniel Bernoulli. Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae [Papers of the Imperial Academy of Sciences in Petersburg]*, 5:175–192, 1738.
- [4] Daniel Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36, January 1954.
- [5] 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.
- [6] Wei Chen, Pinyan Lu, Xiaorui Sun, Bo Tang, Yajun Wang, and Zeyuan Allen Zhu. Optimal pricing in social networks with incomplete information. In *WINE '11: The 7th Workshop on Internet & Network Economics*. Springer, 2011.
- [7] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Knightian truly combinatorial auctions. Technical report, MIT, November 2011.
- [8] Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu. Mechanism design with approximate player types. Technical Report MIT-CSAIL-TR-2011-009, MIT, February 2011.
- [9] 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.
- [10] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [11] Eric Danan. Randomization vs. selection: How to choose in the absence of preference? *Management Science*, 56:503–518, March 2010.

- [12] 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.
- [13] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*, STOC '11, pages 149–158, 2011.
- [14] Uriel Feige and Moshe Tennenholtz. Mechanism design with uncertain inputs: (to err is human, to forgive divine). In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, STOC '11, pages 549–558, New York, NY, USA, 2011. ACM. arXiv:1103.2520v1.
- [15] 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.
- [16] 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.
- [17] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, April 1989.
- [18] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
- [19] Matthew O Jackson. Implementation in undominated strategies: A look at bounded mechanisms. *Review of Economic Studies*, 59(4):757–75, October 1992.
- [20] Frank H. Knight. *Risk, Uncertainty and Profit*. Houghton Mifflin, 1921.
- [21] Giuseppe Lopomo, Luca Rigatti, and Chris Shannon. Uncertainty in mechanism design, 2009. <http://www.pitt.edu/~luca/Papers/mechanismdesign.pdf>.
- [22] Andrew Mas-Colell. An equilibrium existence theorem without complete or transitive preferences. *Journal of Mathematical Economics*, 1(3):237–246, December 1974.
- [23] Paul Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22, Summer 1989.
- [24] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [25] Leandro Nascimento. Remarks on the consumer problem under incomplete preferences. *Theory and Decision*, 70(1):95–110, January 2011.
- [26] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

- [27] Efe A. Ok. Utility representation of an incomplete preference relation. *Journal of Economic Theory*, 104:429–449, 2002.
- [28] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172:1783–1799, October 2008.
- [29] Luca Rigotti and Chris Shannon. Uncertainty and risk in financial markets. *Econometrica*, 73(1):203–243, 01 2005.
- [30] Tuomas Sandholm. Issues in computational vickrey auctions. *International Journal of Electronic Commerce*, 4:107–129, March 2000.
- [31] David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–87, May 1989.
- [32] Wayne Shafer and Hugo Sonnenschein. Equilibrium in abstract economies without ordered preferences. *Journal of Mathematical Economics*, 2(3):345–348, December 1975.
- [33] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.