

Tight Revenue Bounds with Possibilistic Beliefs and Level- k Rationality*

Jing Chen[†] Silvio Micali[‡] Rafael Pass[§]

Abstract

Mechanism design enables a social planner to obtain a desired outcome by leveraging the players' rationality and their beliefs. It is thus a fundamental, but yet unproven, intuition that *the higher the level of rationality of the players, the better the set of obtainable outcomes.*

In this paper we prove this fundamental intuition for players with *possibilistic beliefs*, a model long considered in epistemic game theory. Specifically,

- We define a sequence of *monotonically increasing* revenue benchmarks for single-good auctions, $G^0 \leq G^1 \leq G^2 \leq \dots$, where each G^i is defined over the players' beliefs and G^0 is the second-highest valuation (i.e., the revenue benchmark achieved by the second-price mechanism).
- We (1) construct a single, interim individually rational, auction mechanism that, *without any clue* about the rationality level of the players, guarantees revenue G^k if all players have rationality levels $\geq k + 1$, and (2) prove that no such mechanism can guarantee revenue even close to G^k when at least two players are at most level- k rational.

Keywords: epistemic game theory, incomplete information, single-good auctions

*The authors thank Shafi Goldwasser, Aviad Heifetz, Andrew Lo, Ron Rivest, the editor of *Econometrica* and several wonderful anonymous referees for many helpful comments, and in particular one referee for discussions now reported in Section 5.2. The third author thanks Joseph Halpern for introducing him to the area of epistemic game theory, and for hours and hours of enlightening discussions about it. The first two authors have been partially supported by ONR Grant No. N00014-09-1-0597. The third author has been partially supported by an Alfred P. Sloan Fellowship, Microsoft New Faculty Fellowship, NSF Award CNS-1217821, NSF CAREER Award CCF-0746990, NSF Award CCF-1214844, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2-0211.

[†]Department of Computer Science, Stony Brook University, Stony Brook, NY 11794, USA. jingchen@cs.stonybrook.edu.

[‡]CSAIL, MIT, Cambridge, MA 02139, USA. silvio@csail.mit.edu.

[§]Department of Computer Science, Cornell University, Ithaca, NY 14853, USA. rafael@cs.cornell.edu.

1 Introduction

Mechanism design traditionally models beliefs as probability distributions, and the players as expected-utility maximizers. By contrast, epistemic game theory has successfully and meaningfully studied possibilistic (i.e., set-theoretic) beliefs and more nuanced notions of rationality. In this paper we embrace the possibilistic model and prove that, in single-good auctions, “*more revenue is obtainable from more rational players.*” Let us explain.

Possibilistic (Payoff-type) Beliefs Intuitively, for a player i :

- i 's level-0 beliefs consist of his own (payoff) type;
- i 's level-1 beliefs consist of the *set* of all type subprofiles of his opponents that he considers possible (although he may not be able to compare their relative likelihood);
- i 's level-2 beliefs consist of the set of level-1 belief subprofiles of his opponents that he considers possible;
- and so on.

As usual, beliefs can be wrong¹ and beliefs of different players may be inconsistent; furthermore, we do not assume the existence of a common prior, or that a designer has information about the players' beliefs.

Rationality Following Aumann [9], we do not assume that the players are expected utility maximizers, and allow them to choose actions that are “rational in a minimal sense”. Intuitively,

- A player is (*level-1*) *rational* if he only plays actions that are not *strictly* dominated by some fixed pure action in *every* world he considers possible.²
- Recursively, a player is *level-($k+1$) rational* if he (a) is rational and (b) believes that all his opponents are level- k rational.

We do not assume that a mechanism (designer) has any information about the players' rationality level.

¹That is, a player's belief —unlike his *knowledge*— need not include the true state of the world.

²Due to this notion of rationality, it is without loss of generality to restrict to possibilistic beliefs. If players had probabilistic beliefs, the support of these beliefs alone determines whether a player is rational or not.

Intuitive Description of Our Revenue Benchmarks For auctions of a single good, we consider a sequence of demanding revenue benchmarks, G^0, G^1, \dots .

Intuitively, for any non-negative value v ,

- $G^0 \geq v$ if and only if there exist at least two players valuing the good at least v .
- $G^1 \geq v$ if and only if there exist at least two players believing that there exists a player (whose identity need not be known) valuing the good at least v .
- $G^2 \geq v$ if and only if there exist at least two players believing that there exists a player (whose identity need not be known) believing that there exists a player (whose identity need not be known) valuing the good at least v .
- And so on.

As an example, consider two players, valuing the good 0 and with the following beliefs.

Player 1 believes that player 2

- (a) values the good 100 and
- (b) believes that player 1 values it 200.

Player 2 believes that player 1

- (a') values the good 100 and
- (b') believes that player 2 values it 300.

Then $G^0 = 0$, $G^1 = 100$, and $G^2 = 200$.

It is intuitive (and easily verifiable from the formal definitions) that

- (i) G^0 coincides with the second-highest valuation;
- (ii) $G^0 \leq G^1 \leq \dots$, and each G^{k+1} can be arbitrarily higher than G^k ;
- (iii) If the players' beliefs are correct,³ then each G^k is less than or equal to the highest valuation (but even G^1 can coincide with this valuation);
- (iv) If the players' beliefs are wrong, then even G^1 can be arbitrarily higher than the highest valuation.

³That is, each player considers the true state of the world possible.

Our Results We prove that each additional level of rationality enables one to guarantee a stronger revenue benchmark. Intuitively,

Theorem 1 proves the existence of a *single*, interim individually rational mechanism M that, for all k and all $\varepsilon > 0$, guarantees revenue $\geq G^k - \varepsilon$ whenever the players are *at least* level- $(k + 1)$ rational; and

Theorem 2 proves that, for any k and any $\delta > 0$, no interim individually rational mechanism can guarantee revenue $\geq G^k - \delta$ if at least 2 players are *at most* level- k rational.

A mechanism is *interim individually rational* if each player i , given his true value, has an action guaranteeing him non-negative utility no matter what his opponents might do. (See Subsection 5.2 for further discussion.)

We stress that the guarantee of Theorem 1 is stronger than saying “For each k there exists a mechanism M_k guaranteeing revenue $\geq G^k - \varepsilon$ whenever the players are level- $(k + 1)$ rational.” By contrast, our mechanism M has no information about the players’ rationality levels: it *automatically* guarantees revenue $\geq G^k - \varepsilon$ when the rationality level of each player *happens to be* $\geq k + 1$. That is, M returns revenue

$\geq G^0 - \varepsilon$ if the players happen to be level-1 rational;

$\geq G^1 - \varepsilon$ if the players happen to be level-2 rational;

$\geq G^2 - \varepsilon$ if the players happen to be level-3 rational;

and so on.

This guarantee is somewhat unusual: typically a mechanism is analyzed under only one specific solution concept, and thus under one specific rationality level.

2 Related Work

Ever since Harsanyi [40], the players’ beliefs in settings of incomplete information traditionally use probabilistic representations (see Mertens and Zamir [47], Brandenburger and Dekel [24], and the survey by Siniscalchi [54].)

Beliefs that are not probabilistic and players that do not maximize expected utilities have been considered by Ellsberg [34]. He considers beliefs with ambiguity, but in decision theory. Thus his work does not apply to higher-level beliefs or multi-player games. Higher-level beliefs with ambiguity in multi-player games have been studied by Ahn [4]. His work, however, is not concerned with implementation, and relies on

several common knowledge assumptions about the internal consistency of the players' beliefs. Bodoh-Creed [21] characterizes revenue-maximizing single-good auction mechanisms with ambiguity-averse players, but without considering higher-level beliefs, and using a model quite different from ours.⁴ For more works on ambiguous beliefs, see Bewley [20] and the survey by Gilboa and Marinacci [37].

As we shall see in a moment, our belief model is a set-theoretic version of Harsanyi's type structures. Set-theoretic information has also been studied by Aumann [8], but assuming that a player's information about the "true state of the world" is always correct. Independently, set-theoretic models of beliefs have been considered, in modal logic, by Kripke [46] (see [36] for a well written exposition).

In [26], the first two authors of this paper considered single-good auctions where the players' only have level-1 possibilistic beliefs, and construct a mechanism achieving the benchmark G^1 under a new solution concept, conservative strict implementation. (In particular, the latter notion of implementation assumes that the players are *expected-utility maximizers*. It is easy to see that level-2 rational implementation implies conservative strict implementation, but not vice versa.)

Robust mechanism design, as initiated by Bergemann and Morris [13], is close in spirit to our work, but studies different questions. In particular, it provides additional justification for implementation in dominant strategies. Although defining social choice correspondences over the players' payoff types only (rather than their arbitrary higher-level beliefs), Bergemann and Morris [15] point out that such restricted social choice correspondences cannot represent revenue maximizing allocations.

In [44], Jehiel considers single-good auctions where the players do not know each other's value distribution and only receive certain forms of coarse feedback from past auctions. He shows that under analogy-based expectation equilibrium [43, 45], the designer can generate more revenue than in the optimal auction characterized by [48, 52]. The approach of [44] and that of ours both assume the players have less structured information about each other compared with the standard Bayesian model. But in [44] it is assumed that the players' values are independently distributed, while in our model a player can believe that the other players' values are correlated with each other (and/or with his own value). Also, the epistemic foundation for

⁴In his model, the players have preferences of the Maximin Expected Utility form, the designer has a prior distribution over the players' valuations, the players' beliefs are always correct (i.e., they all consider the designer's prior plausible), actions coincide with valuations, and the solution concepts used are dominant strategy and Bayesian-Nash equilibrium.

the solution concept used in [44] has not been studied, and it would be interesting to understand what is the weakest assumption about the players’ rationality that is sufficient for them to play an analogy-based expectation equilibrium. Moreover, in [44] the mechanisms are ex-post individually rational (that is, a player can decline an offer after seeing the price), while in our model the mechanisms are interim individually rational. It would be interesting to study how much revenue can be generated under level- k rationality using ex-post individually rational mechanisms.

Higher-level rationality and rationalizability have been studied under different models and contexts; see, e.g., [19, 51, 57, 12, 25, 35, 29, 32, 39] from an epistemic game theoretic view and [55, 49, 56, 41, 28, 22, 27, 30] from an experimental view. In a very recent work [16], Bergemann and Morris introduce *belief-free rationalizability* and establish the informational robustness foundations of this and three other solution concepts. As they have pointed out, belief-free rationalizability coincides with our solution concept in private value games —that is, a player’s utility only depends on his own payoff type and not on those of the other players. Different from the above-mentioned studies, our results show that higher-level rationality and rationalizability are also useful solution concepts in mechanism design with possibilistic beliefs.

In the Supplementary Materials of this paper, we characterize level- k rationality by means of iterated deletion of strictly dominated strategies. Iterated deletion has been widely used in mechanism design; see, e.g., [1, 53, 14, 6]. Moreover, it has been considered by many as a good metric for measuring the “level of rationality” of the players, not only because it precisely characterizes higher-level rationality under natural rationality and belief models, but also because of its empirical explanatory power (see [5] for a recent example). However, the literature also questions whether iterated deletion is the best or the only way to measure the players’ (higher-level) rationality. For example, in our mechanism and those of [1] and [14], the designer relies on the players having “exact rationality”: a player prefers outcome O_1 to O_2 as long as his utility under O_1 is strictly larger than under O_2 , no matter how small the difference is. Thus a player deletes strategies “dominated by an amount of order ε ”, and this elimination is done for k times under level- k rationality (and infinitely many times under common knowledge of rationality). Accordingly, the mechanisms are not robust to the exact rationality of the players. However, this criticism applies to many works in robust mechanism design, and one cannot deal with every form of robustness at once. For a stimulating discussion on this topic, see [1, 38, 2].

3 Our Possibilistic Model

Our model is directly presented for single-good auctions, although it generalizes easily to other strategic settings.

An auction is decomposed into two parts: a *context*, describing the set of possible outcomes and the players (including their valuations and their beliefs), and a *mechanism*, describing the actions available to the players and the process leading from actions to outcomes.

We focus on contexts with finitely many types and on deterministic normal-form mechanisms assigning finitely many (pure) actions to each player. Several variants of our model are discussed in Section S3 of the Supplementary Materials.

Contexts A context C consists of four components, $C = (n, V, \mathcal{T}, \tau)$, where

- n is a positive integer, *the number of players*, and $[n] \triangleq \{1, \dots, n\}$ is *the set of players*.
- V is a positive integer, *the valuation bound*.
- \mathcal{T} , the *type structure*, is a tuple of profiles $\mathcal{T} = (T, \Theta, \nu, B)$ where for each player i ,
 - T_i is a finite set, *the set of i 's possible types*;
 - $\Theta_i = \{0, 1, \dots, V\}$ is *the set of i 's possible valuations*;
 - $\nu_i : T_i \rightarrow \Theta_i$ is *i 's valuation function*; and
 - $B_i : T_i \rightarrow 2^{T^{-i}}$ is *i 's belief correspondence*.
- τ , the *true type profile*, is such that $\tau_i \in T_i$ for all i .

Note that \mathcal{T} is a possibilistic version of Harsanyi's type structure [40] without the players' actions. As usual, in a context $C = (n, V, \mathcal{T}, \tau)$ each player i privately knows his own true type τ_i and his beliefs. Player i 's beliefs are *correct* if $\tau_{-i} \in B_i(\tau_i)$. The profile of *true valuations* is $\theta \triangleq (\nu_i(\tau_i))_{i \in [n]}$.

An outcome is a pair (w, P) , where $w \in \{0, 1, \dots, n\}$ is the *winner* and $P \in \mathbb{R}^n$ is the *price profile*. If $w > 0$ then player w gets the good, otherwise the good is unallocated. If $P_i \geq 0$ then player i pays P_i to the seller, otherwise i receives $-P_i$ from the seller. Each player i 's *utility function* u_i is defined as follows: for each valuation $v \in \Theta_i$ and each outcome (w, P) , $u_i((w, P), v) = v - P_i$ if $w = i$, and $= -P_i$ otherwise. i 's *utility* for an outcome (w, P) is $u_i((w, P), \theta_i)$, more simply $u_i(w, P)$. The *revenue* of outcome (w, P) , denoted by $rev(w, P)$, is $\sum_i P_i$. The set of all contexts

with n players and valuation bound V is denoted by $\mathcal{C}_{n,V}$.

Mechanisms An auction mechanism M for $\mathcal{C}_{n,V}$ specifies

- The set $A \triangleq A_1 \times \cdots \times A_n$, where each A_i is i 's *set of actions*.
We denote the set $\times_{j \neq i} A_j$ by A_{-i} .
- An outcome function, typically denoted by M itself, mapping A to outcomes.

For each context $C \in \mathcal{C}_{n,V}$, we refer to the pair (C, M) as an *auction*.

In an auction, when the mechanism M under consideration is clear, for any player i , valuation v , and action profile a , we may simply use $u_i(a, v)$ to denote $u_i(M(a), v)$, and $u_i(a)$ to denote $u_i(M(a))$.

A mechanism is *interim individually rational (IIR)* if, for every context $C = (n, V, \mathcal{T}, \tau)$ and every player i , there exists some action $a_i \in A_i$ such that for every $a_{-i} \in A_{-i}$,

$$u_i(a) \geq 0.$$

A General Possibilistic Framework In Section S4 of the Supplementary Materials, we present our possibilistic framework for general normal-form games. There, following the principle of epistemic game theory, our goal is to characterize the players' higher-level rationality under possibilistic beliefs rather than to design mechanisms, thus we study a game as a whole instead of decoupling it into a context and a mechanism. Also, since the characterization applies to all possible types of all players, there is no need to specify a true type profile. Theorem S1 characterizes the set of actions consistent with level- k rationality for any integer $k \geq 0$ and Theorem S2 characterizes the set of actions consistent with common-belief of rationality. Below we apply our characterization to auctions.

Rationality By Theorem S1, in a normal-form game with possibilistic beliefs, the notion of (higher-level) rationality of our introduction corresponds to a particular iterative elimination procedure of players' actions. Namely, for every rationality level k , the k -round elimination procedure yields the actions consistent with the players being level- k rational, as follows.

Let $\Gamma = ((n, V, \mathcal{T}, \tau), M)$ be a single-good auction, where $\mathcal{T} = (T, \Theta, \nu, B)$. For each player i , each type $t_i \in T_i$ and each $k \geq 0$, we inductively define $RAT_i^k(t_i)$,

the set of actions consistent with level- k rationality for t_i , or equivalently, the set of level- k rationalizable actions for t_i , in the following manner:

- $RAT_i^0(t_i) = A_i$.
- For each $k \geq 1$ and each $a_i \in RAT_i^{k-1}(t_i)$, $a_i \in RAT_i^k(t_i)$ if there does not exist an alternative action $a'_i \in A_i$ such that $\forall t_{-i} \in B_i(t_i)$ and $\forall a_{-i} \in RAT_{-i}^{k-1}(t_{-i})$,

$$u_i((a'_i, a_{-i}), \nu_i(t_i)) > u_i((a_i, a_{-i}), \nu_i(t_i)),$$

where $RAT_{-i}^{k-1}(t_{-i}) = \times_{j \neq i} RAT_j^{k-1}(t_j)$.

The set of action profiles consistent with level- k rationality for auction Γ is $RAT^k(\tau) \triangleq \times_i RAT_i^k(\tau_i)$.

Remark In a recent work [16], Bergemann and Morris introduce the notion of *belief-free rationalizability* and establish its informational robustness foundation. As they have pointed out, this solution concept coincides with ours in private value games, where a player's utility only depends on his own payoff type and not on those of the other players. Intuitively, our solution concept is based on elimination of strictly dominated strategies, and belief-free rationalizability is based on elimination of never-best responses to beliefs about other players' correlated strategies and types. The equivalence of the two solution concepts in private value games follows from the fact that being strictly dominated is the same as being a never-best response to beliefs about correlated strategies and types (see, e.g., [50]). In games with interdependent types, the two notions are different and belief-free rationalizability implies our solution concept.

Level- k Rational Implementation A revenue benchmark b is a function mapping contexts to reals.

Definition 1. A mechanism M level- k rationally implements a revenue benchmark b for $\mathcal{C}_{n,V}$ if, for every context $C \in \mathcal{C}_{n,V}$ and every profile a that is consistent with level- k rationality for auction (C, M) ,

$$rev(M(a)) \geq b(C).$$

Notice that our notion of implementation does not require that the players have

the same level of rationality. Since $RAT^{k'}(\tau) \subseteq RAT^k(\tau)$ for any $k' \geq k$, if a mechanism level- k rationally implements b , then it guarantees b as long as all players have rationality levels $\geq k$.

Furthermore, our notion of implementation does *not* depend on common belief of rationality (a very strong assumption); does *not* require any consistency about the beliefs of different players; and is by definition “closed under Cartesian product.”⁵

Finally, let us stress that in our notion the mechanism knows only the number of players and the valuation bound. (One may consider weaker notions where the mechanism is assumed to know —say— the entire underlying type structure, but not the players’ true types. Of course more revenue benchmarks might be implementable under such weaker notions.)

4 Our Revenue Benchmarks

Below we recursively define the revenue benchmarks G^k for single-good auctions, based on the players’ level- k beliefs. Each G^k is a function mapping a context $C = (n, V, \mathcal{T}, \tau)$ to a real number. For simplicity we let $\max\{v\} \triangleq \max\{v_1, \dots, v_n\}$ for every profile $v \in \mathbb{R}^n$.

Definition 2. Let $C = (n, V, \mathcal{T}, \tau)$ be a context where $\mathcal{T} = (T, \Theta, \nu, B)$. For each player i and each integer $k \geq 0$, the function g_i^k is defined as follows: $\forall t_i \in T_i$,

$$g_i^0(t_i) = \nu_i(t_i) \quad \text{and} \quad g_i^k(t_i) = \min_{t'_{-i} \in B_i(t_i)} \max\{(g_i^{k-1}(t_i), g_{-i}^{k-1}(t'_{-i}))\} \quad \forall k \geq 1.$$

We refer to $g_i^k(t_i)$ as the *level- k guaranteed value of i with type t_i* .

The *level- k revenue benchmark* G^k maps C to the second highest value in $\{g_i^k(\tau_i)\}_{i \in [n]}$.

For any $\varepsilon > 0$, $G^k - \varepsilon$ is the revenue benchmark mapping every context C to $G^k(C) - \varepsilon$.

Note that, if $g_i^k(t_i) \geq c$, then player i with type t_i believes that there always exists some player $j^{(1)}$ —possibly unknown to i — who believes that there always exists

⁵For a given solution concept S this means that S is of the form $S_1 \times \dots \times S_n$, where each S_i is a subset of i ’s actions. This property is important as it overcomes the “epistemic criticism” of the Nash equilibrium concept, see [11, 10, 7]. Indeed, implementation at all Nash equilibria is not closed under Cartesian product, and thus mismatches in the players’ beliefs (about each other’s equilibrium actions) may easily yield undesired outcomes.

a player $j^{(2)}$... who believes that there always exists some player $j^{(k)}$ whose true valuation is at least c .

In Subsection 5.2, we provide some simple examples that illustrate our benchmarks and how our mechanism works.

Remark Note that the values g_i^k 's are monotonically non-decreasing in k . Indeed,

$$g_i^k(t_i) = \min_{t'_{-i} \in B_i(t_i)} \max\{g_i^{k-1}(t_i), g_{-i}^{k-1}(t'_{-i})\} \geq \min_{t'_{-i} \in B_i(t_i)} g_i^{k-1}(t_i) = g_i^{k-1}(t_i).$$

Thus $G^k(C) \geq G^{k-1}(C)$ for every context C and $k > 0$. $G^0(C)$ is the second highest true valuation. It is easy to see that, for every context C , if the players' beliefs are correct, then for each player i and each $k \geq 0$, we have $g_i^k(\tau_i) \leq \max_j \theta_j$, and thus $G^k(C) \leq \max_j \theta_j$.

5 Tight Revenue Bounds

5.1 The Mechanism and the Lower Bounds

While the players' beliefs may be arbitrarily complex, we now show that they can be successfully leveraged by a normal-form mechanism that asks the players to report very little information. Roughly speaking, our mechanism pays the players to receive information about their beliefs, and then uses such information to set a sophisticated reserve price in an otherwise ordinary second-price auction. The idea of buying information from the players is not new; see, e.g., [31] and [33]. There is also a literature in mechanism design that investigates the possibility of "buying higher-level beliefs", such as [3, 14, 18, 26, 17]. However, such studies consider settings with a common prior or settings of complete information, or focus on buying first-level beliefs only. We are not aware of any mechanism where arbitrary higher-level beliefs up to arbitrary levels are being bought without assuming a common prior or complete information.

Our mechanism does not just pay to receive *information* about the players beliefs. It pays to hear even the *faintest rumors* about them. A bit more precisely, it elicits information about the players' beliefs up to some *level bound* K that can be arbitrarily high. For example, if $K = 99$, then (without any information about the

players' rationality level) the mechanism incentivizes a player whose rationality level happens to be $k + 1$, where $k \leq 99$, to report information about his beliefs up to level k . However, the mechanism does not provide incentives for the players to report information about beliefs whose level is greater than 99.

Our mechanism is uniformly constructed on parameters n , V , K , and $\varepsilon > 0$.⁶ An action of a player i has three components: his own identity (for convenience only), a *belief-level* $\ell_i \in \{0, 1, \dots, K\}$, and a *value* $v_i \in \{0, 1, \dots, V\}$. In the description below, the players act only in Step **1**, and Steps **a** through **c** are just “conceptual steps taken by the mechanism”. The expression “ $X := x$ ” denotes the operation that sets or resets variable X to value x .

Mechanism $M_{n,V,K,\varepsilon}$

- 1:** *Each player i , publicly and simultaneously with the others, announces a triple $(i, \ell_i, v_i) \in \{i\} \times \{0, 1, \dots, K\} \times \{0, 1, \dots, V\}$.*
- a:** *Order the n announced triples according to v_1, \dots, v_n decreasingly, and break ties according to ℓ_1, \dots, ℓ_n increasingly. If there are still ties, then break them according to the players' identities increasingly.*
- b:** *Let w be the player in the first triple, $P_w := 2^{nd}v \triangleq \max_{j \neq w} v_j$, and $P_i := 0 \forall i \neq w$.*
- c:** *$\forall i, P_i := P_i - \delta_i$, where $\delta_i \triangleq \frac{\varepsilon}{2n} \left[1 + \frac{v_i}{1+v_i} - \frac{\ell_i}{(1+\ell_i)(1+v_i)^2} \right]$.*

The final outcome is (w, P) . We refer to δ_i as player i 's reward.

Note that our mechanism never leaves the good unsold.

Remark Allegedly, if i is level- k rational, then $v_i = g_i^{k-1}(\tau_i)$ and $\ell_i = \min\{\ell : g_i^\ell(\tau_i) = g_i^{k-1}(\tau_i)\}$. That is, v_i is the highest value v such that i believes “*there exists some player who believes*” ... ($k - 1$ times) some player values the good v , and ℓ_i is the smallest level of beliefs about beliefs needed to attain v_i . Roughly speaking, v_i is the highest “rumored” valuation according to player i 's level- $(k - 1)$ beliefs, and

⁶The reliance on V and K is only to ensure that our mechanism has a finite action space, because our characterization of level- k rationality is for finite games: see Section S3 for more discussion on finite v.s. infinite action spaces.

ℓ_i is the “closeness” of the rumor.⁷ We would like to emphasize that, following the definition of our benchmark, we only require player i to believe “*there exists some player who believes ...*”, instead of “*all players believe ...*”: player i only reports a rumor he believes true for somebody (whose identity he may not even know), rather than a rumor he believes true for everybody. If we were to require the latter, then the benchmark $g_i^k(t_i)$ would have been defined as $\min_{t'_{-i} \in B_i(t_i)} \min\{(g_i^{k-1}(t_i), g_{-i}^{k-1}(t'_{-i}))\}$.

We have the following theorem.

Theorem 1. *For any n, V, K and $\varepsilon > 0$, the mechanism $M_{n,V,K,\varepsilon}$ is IIR and, for each $k \in \{0, 1, \dots, K\}$, level- $(k+1)$ rationally implements the benchmark $G^k - \varepsilon$ for $\mathcal{C}_{n,V}$.*

Note that $M_{n,V,K,\varepsilon}$ does *not* depend on k and is *not* told what the players’ rationality level is. Rather, $M_{n,V,K,\varepsilon}$ *automatically* produces revenue $G^k - \varepsilon$ in every play in which the players *happen to be* level- $(k+1)$ rational. Indeed, (1) such players use only actions that are consistent with level- $(k+1)$ rationality and, (2) at each profile a of such actions (as per Definition 1), $rev(M_{n,V,K,\varepsilon}(a)) \geq G^k - \varepsilon$.

Theorem 1 is proved in Section S1 of the Supplementary Material. Below we provide a proof sketch that highlights the key ideas. In Section S1 and in the discussion below, we more simply denote $M_{n,V,K,\varepsilon}$ by M .

Proof sketch. To prove our revenue lower bound, the key is to prove that a player i “does not underbid”: that is, for any $k \geq 1$ and $a_i = (i, \ell_i, v_i) \in RAT_i^k(\tau_i)$, either $v_i > g_i^{k-1}(\tau_i)$, or $v_i = g_i^{k-1}(\tau_i)$ and $\ell_i \leq \min\{\ell : g_i^\ell(\tau_i) = g_i^{k-1}(\tau_i)\}$. Notice that the way “no underbidding” is defined is consistent with the way how M breaks ties: that is, if player i underbids, his rank after Step **a** of M can only get worse.

As part of the reason for “no underbidding”, we can show that player i ’s reward δ_i strictly decreases if he underbids. A detailed proof for this fact can be found in Claim 2 of Section S1, and below we simply rely on this fact.

The proof proceeds by induction on k , and in this sketch we focus on the case where $k > 1$ and $\hat{\ell}_i \triangleq \min\{\ell : g_i^\ell(\tau_i) = g_i^{k-1}(\tau_i)\} \geq 1$. Arbitrarily fixing the other players’ type subprofile $t_{-i} \in B_i(\tau_i)$ and action subprofile $a'_{-i} \in RAT_{-i}^{k-1}(t_{-i})$, we want to compare player i ’s utility when he bids $\hat{\ell}_i$ and $\hat{v}_i = g_i^{k-1}(\tau_i)$ to his utility when he

⁷We could have defined the mechanism to break ties lexicographically: that is, it first orders the announced triples according to the v_i ’s decreasingly and then according to the ℓ_i ’s decreasingly as well. All the analysis still holds after changing the definition of δ_i respectively. However, such a lexicographical ordering does not have an intuitive explanation for the ℓ_i ’s as we have discussed.

underbids with some ℓ_i and v_i . The crucial part is to show that, by bidding $\hat{a}_i = (i, \hat{\ell}_i, \hat{v}_i)$ player i does not win the good, thus neither does he by bidding $a_i = (i, \ell_i, v_i)$, which is ranked even worse. Notice why this would conclude the proof: since player i does not win the good in either case, his utilities are exactly the rewards he gets under \hat{a}_i and a_i respectively. Since the reward strictly decreases when he underbids, his utility is strictly smaller by bidding a_i , which (together with the analysis of several other cases) implies that a_i is strictly dominated by \hat{a}_i and thus cannot be level- k rationalizable.

To see why player i does not win the good by bidding \hat{a}_i , on the one hand, notice that by the definition of $\hat{\ell}_i$ we have $g_i^{\hat{\ell}_i-1}(\tau_i) < g_i^{\hat{\ell}_i}(\tau_i)$. On the other hand, by the definition of $g_i^{\hat{\ell}_i}(\tau_i)$ there exists a player j such that $g_j^{\hat{\ell}_i-1}(t_j) \geq g_i^{\hat{\ell}_i}(\tau_i)$. Accordingly, $j \neq i$. By the inductive hypothesis, player j does not underbid, which means he bids at least $(\hat{\ell}_i - 1, g_j^{\hat{\ell}_i-1}(t_j))$. Thus the mechanism ranks j ahead of i , and i cannot be the winner, as we wanted to show. ■

Remark Very roughly speaking, having $\hat{\ell}_i \geq 1$ means that player i believes that somebody else will bid at least $\hat{\ell}_i - 1$ and $g_i^{\hat{\ell}_i}(\tau_i)$, thus i bids at least $\hat{\ell}_i$ and $g_i^{\hat{\ell}_i}(\tau_i)$, indicating that he does not want to win the good at his own bid, rather he is contributing his beliefs about others' values so as to receive a better reward. In Subsection 5.2 we further elaborate on this phenomenon by means of a few examples and clarify the revenue dependency on the players' rationality and the adopted solution concept.

5.2 Understanding Theorem 1

We start with a setting of complete information.

An Auction with Complete Information Consider the following example. There are two players, Player 1 has a unique type t_1 where her value is 100, player 2 has a unique type t_2 where his value is 200, and the setting is of complete information (that is, $\tau = (t_1, t_2)$, $B_1(t_1) = \{t_2\}$ and $B_2(t_2) = \{t_1\}$). In this example the level-0 benchmark, G^0 , is 100, while any G^k for $k > 0$ is 200. Indeed, player 1 believes that player 2's value is 200 and player 2 knows that his value is 200, etc. Let us now analyze the revenue performance of mechanism M in this setting under different rationality levels of the players.

When the players are (level-1) rational, they use actions that survive one round of elimination. In particular, player 1 will report $\ell_1 = 0$ and $v_1 \geq 100$, while player 2 will report $\ell_2 = 0$ and $v_2 \geq 200$. Indeed, each player i prefers to win the good at a price lower than his/her own valuation and to report the lowest belief level ℓ_i , because his/her reward δ_i increases with reporting higher values and lower belief levels. Thus mechanism M virtually guarantees the revenue benchmark G^0 when the players are level-1 rational.

When the players are level-2 rational (i.e., they have mutual belief of rationality), each one of them believes that the other will use actions that survive one round of elimination and further eliminates his/her own actions based on such beliefs. In particular, player 1 believes that player 2 will report $\ell_2 = 0$ and $v_2 \geq 200$. Therefore player 1 believes that she herself will surely lose by reporting $v_1 = 100$, so she will increase her bid in order to get a bigger reward. However, if she bids $\ell_1 = 0$ and $v_1 \geq 200$, and if player 2 bids exactly $v_2 = 200$, then, according to the tie-breaking rule, she will win the auction at price 200, thus getting a negative utility. By bidding $\ell_1 = 1$ and $v_1 = 200$ she indicates that *she does not want to win but only to get a better reward*, and avoids buying the good at more than what she values it. Indeed, player 2 will continue bidding $\ell_2 = 0$ and $v_2 \geq 200$, winning the good at price 200.

In principle there is a tradeoff between player 1's (a) getting a bigger reward by increasing v_1 and (b) getting a smaller reward due to the increase of ℓ_1 from 0 to 1. However, the reward function is designed so that increasing both v_i and ℓ_i by one unit actually increases the reward. This is the sense in which player 1 is rewarded for reporting the "rumor" that player 2's value is 200.

In sum, in this example $G^0 = 100$ and $G^1 = 200$ and, in accordance with Theorem 1, M virtually achieves the revenue benchmark G^0 when the players' rationality level is 1, and the benchmark G^1 when their rationality level is 2. However, $G^1 = G^2 = \dots$. Indeed, in a setting of complete information, higher-level beliefs "collapse" to level-1, all our higher-level benchmarks collapse down to G^1 , and thus M cannot guarantee revenue greater than G^1 .

Our next examples show that, in auctions of *incomplete* information, (1) it is possible that $G^0 < G^1 < G^2 < \dots$ and (2) mechanism M can guarantee revenue virtually equal to G^k when the players' rationality level is $k + 1$, *whether or not all players' beliefs are correct*.

Higher Revenue from Higher Rationality Consider using mechanism M in an auction in which the type structure is as in Figure 1. Here a node represents a type or a type profile, together with the corresponding values of the players; and an edge labeled by $i \in \{1, 2\}$ points to a world that player i believes possible. The subscriptions of the types indicate to which player they belong. For example, the edge from (t_1, t_2) to t_2^1 labeled by 1 means $t_2^1 \in B_1(t_1)$.




Figure 1: A type structure with correct beliefs.

As shown in Figure 1,

- At the true type profile $\tau = (t_1, t_2)$, player 1 values the good 0 and player 2 values it 100; player 2 believes that the only possible world is τ ; and player 1 believes that τ is a possible world, but so is the first (alternative) world t_2^1 , in which player 2 values the good 1.
- In that first world, player 2 believes that the possible world is the second one t_1^2 , where player 1 values the good 2. In general, in the k th world t_i^k with $k < 100$, player i values the good k and believes that the possible world is t_{-i}^{k+1} , where player $-i$ values the good $k + 1$.
- Finally, the 100th world is a setting of complete information where each player values the good 100.

Note that this example is one of incomplete information in which the players hold correct beliefs.⁸ Consider what happens when the players are, say, level-3 rational. By definition, player 1 believes that player 2 is level-2 rational. Also she believes that t_2^1 is a possible type for player 2. What will player 2 do under type t_2^1 when he is

⁸Indeed, having correct beliefs at the true type profile τ only means that all players believe that τ is a possible world, and does not preclude a player from believing that the others may have incorrect beliefs. Thus holding correct beliefs is a much weaker condition than their having complete information at τ . In the latter case, every player believes that τ is the *only* possibility (thus has correct beliefs), every player believes that the others have correct beliefs, and so on.

level-2 rational? Since he believes that player 1 is level-1 rational and has type t_1^2 , he needs to find out what such a player 1 will do. Notice that player 1 has value 2 at t_1^2 . Similar to the example of complete information, when player 1 is level-1 rational and has value 2, she will bid $(0, 2)$. Accordingly, player 2 will bid $(1, 2)$ at t_2^1 , indicating that he believes the good should be sold to player 1 at price 2. Thus, going back to τ , player 1 will bid $(2, 2)$, indicating that she believes that player 2 will bid $(1, 2)$. Since player 2 has true value 100 at τ , he will actually bid $(0, 100)$, and the mechanism sells the good to him at price 2, which is player 1's bid. Roughly speaking, if the players are:

- level- k rational with $1 \leq k \leq 100$, then M gets bids $(k - 1, k - 1)$ and $(0, 100)$, and sells the good for $k - 1$.
- level-101 rational, then M gets bids $(100, 100)$ and $(0, 100)$, and sells the good for 100.

In sum, even when players' beliefs are correct, M can generate more and more revenue as the players' rationality level increases.

Also note that the players having correct beliefs does not imply that they have common knowledge of correct beliefs.⁹ In the latter case, in fact, it is easy to see that $G^k = G^1$ for every $k > 0$. Common knowledge of correct beliefs, however, is a very strong restriction: in particular, standard characterizations of rationalizability [23, 57] do not apply under such restriction [39].

As mentioned before, no G^k can exceed the highest valuation *if the players' beliefs are correct*. However, a simple variant of the current example shows that M can generate arbitrarily high revenue if some player has a suitable incorrect belief. Indeed, consider a type structure that is almost the same as the one in Figure 1, except that at the true type profile τ , player 2 has value 0 and player 1 believes that t_2^1 is the only possible world. In this case, M generates revenue 99 under level-101 rationality, although the highest valuation in the true world is 0.

IIR and Right Reasoning with Wrong Beliefs Notice that our notion of interim individual rationality is stronger than the traditional one in a Bayesian setting. Indeed, the latter notion only requires that a player's *expected* utility (rather than his "actual" utility) is non-negative when his beliefs are correct.

⁹In the current example, player 1 considers it possible that player 2 has incorrect beliefs.

In mechanism M , when a player i 's possibilistic beliefs are *correct*, for any rationality level k he always has a “level- k safe” (or, level- k interim individually rational) action that is consistent with level- k rationality and gives him positive utility against all other players' actions consistent with level- $(k - 1)$ rationality —that is, $v_i = g_i^{k-1}(\tau_i)$ and $\ell_i = \min\{\ell : g_i^\ell(\tau_i) = g_i^{k-1}(\tau_i)\}$.

However, when a player's beliefs are wrong, every action of his that is consistent with level- k rationality may give him negative utility against some actions of the others consistent with level- $(k - 1)$ rationality. For example, consider the case where player 1 has value 0, player 2 has value 100, player 1 believes that player 2 has value 200, and player 2 believes that player 1's true type is the only possible type of player 1. When $k \geq 3$, under any action consistent with level- k rationality player 1 bids at least $(1, 200)$; while a particular action of player 2 that is consistent with level- $(k - 1)$ rationality is $(2, 200)$, which gives player 1 negative utility. Nonetheless, a player still has the same level- k safe action as defined above, and he *believes* that his utility will be non-negative under this action against all other players' actions consistent with level- $(k - 1)$ rationality. This situation is not too dissimilar from that of a rational player who willingly enters the stock market, yet might end up losing money if his beliefs are wrong.¹⁰

Simple Bidding Following our discussion and the analysis in the Supplementary Material, for any $k \geq 0$ and any player i , the structure of actions (i, ℓ, v) consistent with level- k rationality is very easy to describe. In particular, letting $v_i = g_i^{k-1}(\tau_i)$ and $\ell_i = \min\{\ell : g_i^\ell(\tau_i) = g_i^{k-1}(\tau_i)\}$ as before, we have that (1) $v \geq v_i$, and (2) if $v = v_i$ then $\ell \leq \ell_i$. Any action that does not satisfy (1) or (2) is *not* consistent with level- k rationality.

Our mechanism can also be viewed as asking each player to bid just his guaranteed value v_i and then using an endogenous tie-breaking rule which is based on the players' types (including their beliefs); see, e.g., [42] on endogenous tie-breaking rules in games of incomplete information. As pointed out in [42], to implement an endogenous tie-breaking rule (through an actual mechanism) the players are required to announce their whole type. In contrast, in our mechanism each player is only asked to announce (i, ℓ_i, v_i) , which is significantly simpler than announcing the whole type.

¹⁰This exploitation by M of the players' wrong beliefs may appear somewhat unfair. However, we should keep in mind that, when constructing a revenue generating mechanism, a designer works for the seller, and would not do his job properly if he leaves some money on the table.

Variants In Section S3 of the Supplementary Material, we discuss variants of our mechanism M (e.g., dealing with continuous valuation spaces), as well as analyses under different solution concepts.

5.3 The Upper Bounds on Revenue

We now show that level- $(k+1)$ rationality is necessary to guarantee the benchmark G^k .

Theorem 2. *For every n, V, k , and $c < V$, no IIR mechanism level- k rationally implements $G^k - c$ for $\mathcal{C}_{n,V}$ (even if only two players are level- k rational and all others' rationality levels are arbitrarily higher than k).*

Theorem 2 is proved in Section S2 of our Supplementary Materials. Below we provide a proof sketch that explains intuitively why any IIR mechanism that tries to level- k rationally implements $G^k - c$ for $\mathcal{C}_{n,V}$ is bound to fail.

Proof sketch. To see the main ideas it suffices to consider $n = 2$ and $k \geq 1$. Assume there exists an IIR mechanism \hat{M} that level- k rationally implements $G^k - c$ for $\mathcal{C}_{n,V}$. We construct a context $C = (2, V, \mathcal{T}, \tau)$ and show that there exists a level- k rationalizable action profile a such that $rev(\hat{M}(a)) < G^k(C) - c$, contradicting the hypothesis. The type structure $\mathcal{T} = (T, \Theta, \nu, B)$ is defined as follows: for each player i ,

- $T_i = \{t_{i,\ell} : \ell \in \{0, 1, \dots, k\}\}$;
- $\nu_i(t_{i,\ell}) = 0 \forall \ell < k$, and $\nu_i(t_{i,k}) = V$; and
- $B_i(t_{i,\ell}) = \{t_{-i,\ell+1}\} \forall \ell < k$, and $B_i(t_{i,k}) = \{t_{-i,k}\}$.

The type structure \mathcal{T} is illustrated in Figure 2 below, which is the same as Figure 1 in Section S2. We set $\tau_i = t_{i,0}$ for each i .




Figure 2: Type structure \mathcal{T} in context C

To derive the desired contradiction we use an auxiliary context $C' = (2, V, \mathcal{T}', \tau')$, with the type structure $\mathcal{T}' = (T', \Theta, \nu', B')$ defined as follows: for each player i ,

- $T'_i = \{t'_{i,\ell} : \ell \in \{0, 1, \dots, k\}\}$;
- $\nu'_i(t'_{i,\ell}) = 0 \forall \ell$; and
- $B'_i(t'_{i,\ell}) = \{t'_{-i,\ell+1}\} \forall \ell < k$, and $B'_i(t'_{i,k}) = \{t'_{-i,k}\}$.

The type structure \mathcal{T}' is illustrated in Figure 3 below, which is the same as Figure 2 in Section S2. We set $\tau'_i = t'_{i,0}$ for each i .




Figure 3: Type structure \mathcal{T}' in context C'

By induction, we can show that $g_i^k(t_{i,0}) = V$ for each i , thus $G^k(C) = V$ and $G^k(C) - c > 0$. Also by induction, we can show that $RAT^k(\tau) = RAT^k(\tau')$, thus it suffices to show there exists $a \in RAT^k(\tau')$ such that $rev(\hat{M}(a)) \leq 0$. Since both players have value 0 at τ' , it suffices to show that there exists $a \in RAT^k(\tau')$ such that $u_i(a, 0) \geq 0$ for each player i .

To see why this is true, notice that \hat{M} is IIR, which implies that for each i there exists a_i such that $u_i((a_i, a_{-i}), 0) \geq 0$ for any a_{-i} in A_{-i} —the full action set of player i in \hat{M} . Since $B'_i(t'_{i,0}) = \{t'_{-i,1}\}$, by induction we can show that there exists $a_i^k \in RAT_i^k(t'_{i,0})$ such that $u_i((a_i^k, a'_{-i}), 0) \geq 0$ for any $a'_{-i} \in RAT_{-i}^{k-1}(t'_{-i,1})$. We will be done as long as we can show that $a_i^k \in RAT_i^{k-1}(t'_{i,1})$ for each i : in this case, we have $a^k \in RAT^k(\tau')$ and $u_i(a^k, 0) \geq 0$ for each i , as desired.

To see why $a_i^k \in RAT_i^{k-1}(t'_{i,1})$, notice that although $t'_{i,0}$ and $t'_{i,1}$ are different types, player i has the same value 0 under both of them. Also, despite of different names of the types, player i 's beliefs are the same at $t'_{i,0}$ and $t'_{i,1}$ up to level $k - 1$, where player i 's belief at $t'_{i,0}$ reaches $t'_{j,k-1}$ and that at $t'_{i,1}$ reaches $t'_{j,k}$, with $j = i$ if k is odd and $j = -i$ otherwise. Since $RAT_j^0(t'_{j,k-1}) = RAT_j^0(t'_{j,k}) = A_j$, again by induction we can show that $RAT_i^{k-1}(t'_{i,0}) = RAT_i^{k-1}(t'_{i,1})$. Since $a_i^k \in RAT_i^k(t'_{i,0}) \subseteq RAT_i^{k-1}(t'_{i,0})$, $a_i^k \in RAT_i^{k-1}(t'_{i,1})$ as we wanted to show, concluding the proof. ■

Remark Roughly speaking, this is what may go wrong if a mechanism tries to level- k rationally implement $G^k - c$: under the type structure \mathcal{T} and true type profile τ ,

the players can *pretend* that their true type profile is τ' , and the mechanism, looking only at level- k rationalizable actions, cannot distinguish these two cases.

Also notice that the players have wrong beliefs in the type structures used in our proof. Type structures with correct beliefs can be used to prove our theorem for some values of c (more specifically, for any $c < V/2$), but in any such type structure some players necessarily believe that the others may have incorrect beliefs. Indeed, the type structures constructed for proving Theorem 2 are necessarily inconsistent with a common prior: as mentioned towards the end of the second example in Section 5.2, when there is a common prior, we have $G^k = G^1$ for all $k \geq 1$, thus our own mechanism level- k rationally implements $G^k - \varepsilon$ for $\mathcal{C}_{n,V}$ for any $k \geq 2$.

6 Conclusions and Future Directions

Mechanism design enables a social planner to obtain a desired outcome by leveraging the players' rationality and beliefs. It is thus a fundamental intuition that “the higher the players' rationality level, the better the obtainable outcomes”. Theorems 1 and 2 prove this intuition under possibilistic beliefs and Aumann's notion of rationality.

Let us remark that our mechanism M of Theorem 1 is also applicable in a Bayesian framework, by letting the players' possibilistic beliefs be the support of their probabilistic ones. However, due to our use of Aumann's notion of rationality, our analysis has a bite only when the supports of these beliefs may be different. (For instance, in the case of “full-support” beliefs, all our benchmarks collapse down to G^1 .)

It would be interesting to extend our analysis to Bayesian settings using expected utility maximization for the underlying notion of rationality. In particular, in such a setting, the safe action for a level- k rational player (which guarantees him non-negative utility when his beliefs are correct) may be dominated by an “overbidding” strategy with positive expected utility that, seeking a higher reward, may sometimes result in receiving negative utility.

Finally, let us remark that, although we provide revenue bounds in a possibilistic setting with higher-level beliefs, we do so without actually identifying an *optimal* mechanism in the sense of Myerson [48]. Identifying such a mechanism is another interesting open problem.

References

- [1] D. Abreu and H. Matsushima. Virtual Implementation in Iteratively Undominated Actions: Complete Information. *Econometrica*, 60(5): 993-1008, 1992.
- [2] D. Abreu and H. Matsushima. A Response to Glazer and Rosenthal. *Econometrica*, Vol. 60, No. 6, pp. 1439-1442, 1992.
- [3] D. Abreu and H. Matsushima. Virtual Implementation in Iteratively Undominated Actions: Incomplete Information. Working paper, 1992.
- [4] D. S. Ahn. Hierarchies of ambiguous beliefs. *Journal of Economic Theory*, Vol. 136, pp. 286-301, 2007.
- [5] L. Alaoui and A. Penta. Endogenous Depth of Reasoning. Under revision at *Review of Economic Studies*, 2014.
- [6] G. Artemov, T. Kunimoto, and R. Serrano. Robust Virtual Implementation with Incomplete Information: Towards a Reinterpretation of the Wilson Doctrine. *Journal of Economic Theory*, 148, 424-447, 2013.
- [7] G. B. Asheim, M. Voorneveld, and J. W. Weibull. Epistemically stable action sets. Working paper, 2009.
- [8] R. Aumann. Agreeing to Disagree. *Annals of Statistics*, 4: 1236-1239, 1976.
- [9] R. Aumann. Backward Induction and Common Knowledge of Rationality. *Games and Economic Behavior*, Vol. 8, pp. 6-19, 1995.
- [10] R. Aumann and A. Brandenburger. Epistemic Conditions for Nash Equilibrium. *Econometrica*, Vol. 63, No. 5, pp. 1161-1180, 1995.
- [11] K. Basu and J.W. Weibull. Action subsets closed under rational behavior. *Economics Letters*, Vol. 36, pp. 141-146, 1991.
- [12] P. Battigalli and M. Siniscalchi. Rationalization and incomplete information. *B. E. Journal of Theoretical Economics*, Vol. 3, Iss. 1, 2003.
- [13] D. Bergemann and S. Morris. Robust mechanism design. *Econometrica*, Vol. 73, No. 6, pp. 1771-1813, 2005.
- [14] D. Bergemann and S. Morris. Robust Virtual Implementation. *Theoretical Economics*, Vol. 4, pp. 45-88, 2009.
- [15] D. Bergemann and S. Morris. Robust Mechanism Design: An Introduction. In D. Bergemann and S. Morris, *Robust Mechanism Design*, World Scientific

Press, 2012.

- [16] D. Bergemann and S. Morris. Informational Robustness and Solution Concepts. Working paper, 2014.
- [17] D. Bergeman, S. Morris, and S. Takahashi. Efficient Auctions and Interdependent Types. *American Economic Review Papers and Proceedings*, 102(3): 319-324, 2012.
- [18] D. Bergemann, S. Morris, and O. Tercieux. Rationalizable Implementation. *Journal of Economic Theory*, 2011.
- [19] B. Bernheim. Rationalizable Strategic Behavior. *Econometrica*, Vol. 52, No. 4, pp. 1007-1028, 1984.
- [20] T. F. Bewley. Knightian decision theory. Part I. *Decisions in Economics and Finance*, Vol. 25, pp. 79-110, 2002.
- [21] A. L. Bodoh-Creed. Ambiguous beliefs and mechanism design. *Games and Economic Behavior*, Vol. 75, pp. 518-537, 2012.
- [22] A. Bosch-Domènech, J. G. Montalvo, R. Nagel, and A. Satorra. One, Two, (Three), Infinity, ...: Newspaper and Lab Beauty-Contest Experiments. *American Economic Review*, 92, 1687-1701, 2002.
- [23] A. Brandenburger and E. Dekel. Rationalizability and correlated equilibria. *Econometrica*. Vol. 55, pp. 1391-1402, 1987.
- [24] A. Brandenburger and E. Dekel. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, Vol. 59, pp. 189-198, 1993.
- [25] C. Camerer, T.-H. Ho, and J.-K. Chong. A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics*, 119, 861-898, 2004.
- [26] J. Chen and S. Micali. Mechanism Design with Set-Theoretic Beliefs. *Symposium on Foundations of Computer Science (FOCS)*, pp. 87-96, 2011. An improved version appears in *Journal of Economic Theory*, Vol. 156, pp. 77-102, 2015.
- [27] M. A. Costa-Gomes and V. P. Crawford. Cognition and Behavior in Two-Person Guessing Games: An Experimental Study. *American Economic Review*, 96, 1737-1768, 2006.
- [28] M. A. Costa-Gomes, V. P. Crawford, and B. Broseta. Cognition and Behavior in Normal-Form Games: An Experimental Study. *Econometrica*, 69: 1193-1235, 2001.

- [29] V. P. Crawford and N. Iriberri. Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner’s Curse and Overbidding in Private-Value Auctions? *Econometrica*, Vol. 75, No. 6, pp. 1721-1770, 2007.
- [30] V. P. Crawford and N. Iriberri. Fatal Attraction: Salience, Naivete, and Sophistication in Experimental ‘Hide-and-Seek’ Games. *American Economic Review*, 97, 1731-1750, 2007.
- [31] J. Crémer and R. P. Mclean. Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent. *Econometrica*, Vol. 53, No. 2, pp. 345-361, March 1985.
- [32] E. Dekel, D. Fudenberg, S. Morris. Interim correlated rationalizability. *Theoretical Economics*, Vol. 2, pp. 15-40, 2007.
- [33] K. Eliaz and R. Spiegler. A Mechanism-Design Approach to Speculative Trade. *Econometrica*, Vol. 75, No. 3, pp. 875-884, 2007.
- [34] D. Ellsberg. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, Vol. 75, pp. 643-669, 1961.
- [35] J. C. Ely and M. Peşki. Hierarchies of belief and interim rationalizability. *Theoretical Economics*, Vol. 1, pp. 19-65, 2006.
- [36] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT Press, 2003.
- [37] I. Gilboa and M. Marinacci. Ambiguity and the Bayesian paradigm. Working paper, 2011.
- [38] J. Glazer and R. Rosenthal. A Note on Abreu-Matsushima Mechanisms. *Econometrica*, Vol. 60, No. 6, pp. 1435-1438, 1992.
- [39] J. Halpern and R. Pass. Conservative belief and rationality. *Games and Economic Behavior*, Vol. 80, pp. 186-192, 2013.
- [40] J. Harsanyi. Games with Incomplete Information Played by “Bayesian” Players, I-III. *Management Science*, Vol. 14, pp. 159-182, 320-334, 486-502. 1967-1968.
- [41] T.-H. Ho, C. Camerer, and K. Weigelt. Iterated Dominance and Iterated Best Response in Experimental p-Beauty Contests. *American Economic Review*, Vol. 88, pp. 947-969, 1998.
- [42] M. Jackson, L. K. Simon, J. M. Swinkels, and W. R. Zame. Communication and Equilibrium in Discontinuous Games of Incomplete Information. *Econometrica*,

- Vol. 70, No. 5, pp. 1711-1740, 2002.
- [43] P. Jehiel. Analogy-based expectation equilibrium. *Journal of Economic Theory*, Vol. 123, pp. 81-104, 2005.
 - [44] P. Jehiel. Manipulative auction design. *Theoretical Economics*, 6: 185-217, 2011.
 - [45] P. Jehiel and F. Koessler. Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior*, 62, 533-557, 2008.
 - [46] S. Kripke. Semantical analysis of modal logic I: normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, Vol. 9, pp. 67-96, 1963.
 - [47] J. Mertens and S. Zamir. Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14: 1-29, 1985.
 - [48] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, Vol.6, No.1, pp. 58-73, 1981.
 - [49] R. Nagel. Unraveling in Guessing Games: An Experimental Study. *American Economic Review*, Vol. 85, pp. 1313-1326, 1995.
 - [50] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*, MIT Press, 1994.
 - [51] D. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, Vol. 52, No. 4, pp. 1029-1050, 1984.
 - [52] J. G. Riley and W. F. Samuelson. Optimal auctions. *American Economic Review*, 71: 381-392, 1981.
 - [53] R. Serrano and R. Vohra. A Characterization of Virtual Bayesian Implementation. *Games and Economic Behavior*, 50, 312-331, 2005.
 - [54] M. Siniscalchi. Epistemic Game Theory: Beliefs and Types. In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*, Second Edition. Palgrave Macmillan, 2008.
 - [55] D. Stahl and P. Wilson. Experimental Evidence on Players' Models of Other Players. *Journal of Economic Behavior and Organization*, 25: 309-327, 1994.
 - [56] D. Stahl and P. Wilson. On Players' Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, Vol. 10, pp. 213-254, 1995.
 - [57] T. Tan and S. Werlang. The Bayesian foundation of solution concepts of games. *Journal of Economic Theory*, Vol. 45, pp. 370-391, 1988.