Foundations of Dominant Strategy Mechanisms

Kim-Sau Chung
Department of Economics
Northwestern University
sau@nwu.edu

Jeffrey C. Ely*
Department of Economics
Boston University
ely@bu.edu

January 23, 2004

Abstract

Wilson (1987) criticizes the existing literature of game theory as relying too much on common-knowledge assumptions. In reaction to Wilson's critique, the recent literature of mechanism design has started to employ stronger solution concepts such as dominant strategy incentive compatibility, and restrict attention to simpler mechanisms such as dominant strategy mechanisms. However, there has been little theory behind this approach. In particular, it has not been made clear why employing simpler mechanisms, instead of more complicated ones, is the correct way to address Wilson's critique. This paper aims at filling this void. We propose a potential theory, known as the *maxmin* theory, which postulates that a *cautious* mechanism designer, facing uncertainty over which (common-knowledge) assumptions are valid and which are not, would indeed rationally choose simpler mechanisms such as dominant strategy mechanisms. In this paper, we summarize our progress in proving this theory, explore other possible theories, and discuss related theoretical questions that will be of interest in other areas.

*Support from the National Science Foundation under grant #SES 99-85462 is gratefully acknowledged.

1 Introduction

In the recent literature of mechanism design, there is a research agenda which is motivated by the so-called *Wilson Doctrine*. Roughly speaking, the Wilson Doctrine refers to Wilson's (1987) vision that a good theory of mechanism design should not rely too heavily on assumptions of common knowledge:

> "Game theory has a great advantage in explicitly analying the consequences of trading rules that presumably are really common knowledge; it is deficient to the extent it assumes other features to be common knowledge, such as one agent's probability assessment about another's preferences or information. [...] I foresee the progress of game theory as depending on succesive reduction in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality."

Although there is no clear prescription from Wilson (1987) on how exactly to achieve such an ideal of "successive reduction" in the dependence on common knowledge assumptions, recent researches have converged to the idea that the appropriate way to implement the Wilson Doctrine is to use stronger solution concepts. For example, when Dasgupta and Maskin (2000) and Perry and Reny (2002) design efficient auctions in interdependent-value settings, they insist that their designs are ex post incentive compatible. Similarly, when Segal (2002) designs optimal auctions in private-value settings, he also insists that his designs are dominant strategy incentive compatible. Both ex post incentive compatibility and dominant strategy incentive compatibility are stronger solution concepts compared with Bayesian incentive compatibility, which in turn is the solution concept used in the traditional literature on mechanism design. In private-value settings, ex post incentive compatibility boils down to dominant strategy incentive compatibility.

In this paper, we ask the question of in what sense using these stronger solution concepts is an appropriate way to implement Wilson Doctrine. In particular, we focus on private-value settings, and ask whether or not there is a foundation for using dominant strategy mechanisms.

There are two reasons why this is an important question. At a shallower level, it is easier to be destructive but more difficult to be constructive. It is easier to dismiss the previous literature as "deficient," but more difficult to establish a particular new approach as *the* correct approach among many other possible alternatives. Whenever a particular new approach is suggested, we should insist that there is a theory behind the suggestion.

At a deeper level, it is not apparent at all why, when a mechanism design is not willing to make strong common knowledge assumptions, she would then use simpler mechanisms such as dominant strategy mechanisms, instead of using mechanisms that are even more complicated. In principle, a mechanism designer can ask her agents *anything* that she does not know, and she *should* do so if the answers are potentially useful. For example, if she is not sure whether a certain common knowledge assumption is true or not, she can (and probably *should*) add to her original mechanism an additional question concerning the validity of this common knowledge assumption. The fewer assumptions the mechanism designer is willing to make, the more questions she should ask, and hence the more complicated her mechanism should be. Pushing this logic to its extreme, if we were ever to achieve Wilson's ideal of "successive reduction" in the dependence of common knowledge assumptions, we would envision mechanisms that are so complicated that they ask agents to report *everything*. At the limit, mechanisms would become so complicated that they ask agents to report their whole infinite hierarchies of beliefs and higher-order beliefs, or in other words to report their universal types. It seems that the suggestion of using simpler mechanisms such as dominant strategy mechanisms is squarely at odd with this established intuition in the literature of mechanism design.

In this paper, we shall provide a rationale for using dominant strategy mechanisms that would also reconcile the above intuition. Our theory can be loosely explained with the following story. Imagine the mechanism designer as an auctioneer. She may have confidence in her estimate of the distribution ν of the bidders' valuations, perhaps based on data from similar auctions in the past. But she does not have reliable information about the bidders' beliefs (including their beliefs about one another's valuations, their beliefs about these beliefs, etc.), as these are arguably never observed. She can choose to use a dominant strategy mechanism, or she can alternatively choose to use some Bayesian incentive compatible mechanism that will perform much better under certain common knowledge assumptions. It is well known that the performance (which means the auctioneer's profit in this auction example) of a dominant strategy mechanism is insensitive to the bidders' beliefs, on which the auctioneer does not have reliable information anyway. On the other hand, a Bayesian incentive compatible mechanism that performs well under certain common knowledge assumptions may perform horribly if those assumptions turn out to be false. Thus, the auctioneer, faced with uncertainty about the bidders' beliefs, may optimally choose to play safe and use a dominant strategy mechanism.

Note that this story by itself does not "assume away" the availability of complicated mechanisms that ask bidders to report their universal types. These complicated mechanisms are still available. However, in order to guarantee that bidders will be honest, these mechanisms have to be at least Bayesian incentive compatible. In general, the performance of a Bayesian incentive compatible mechanism sensitively depends on which assumptions (about bidders' beliefs) are true and which are false, unless it is also a dominant strategy mechanism. Therefore a cautious auctioneer's voluntarily chooses to use a dominant strategy mechanism.

We call this story the *maxmin* foundation of dominant strategy mechanisms, because the auctioneer chooses among mechanisms according to their worst scenerio performance. Pictorially, what we need to prove is a theorem that takes the form of Figure 1. In Figure 1, we plot the performance of an arbitrary Bayesian incentive compatible mechanism against different assumptions about (or distributions of) bidders' beliefs. The graph of any dominant strategy mechanism—and in particular the graph of the best one among all dominant strategy mechanisms—will be a horizontal line. What we need to prove is that the graph of any (potentially very complicated) Bayesian incentive compatible mechanism must dip below the graph of the best dominant strategy mechanism at some point.

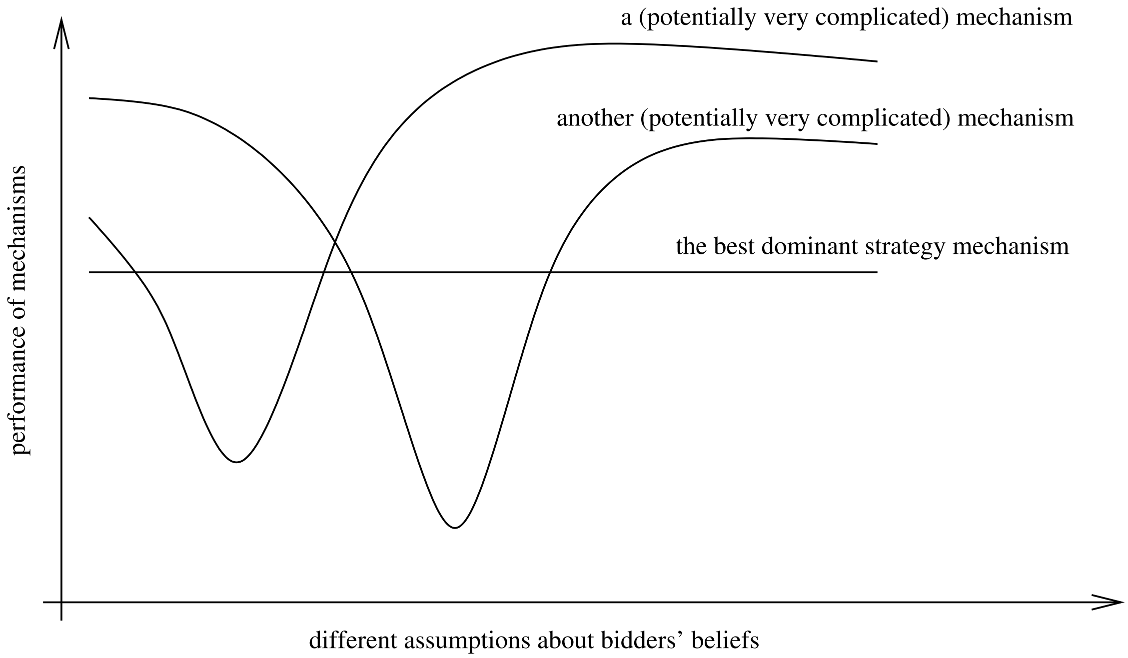

Figure 1: the graph of any mechanism dips below the graph of the best dominant strategy mechanism at some point.

Figure 1, although we believe captures the imagination of many advocates of dominant strategy mechanisms, turns out being very difficult to prove. The set of mechanisms is simply too rich, especially when we allow for complicated mechanisms, and at this stage we still do not know how to systematically characterize the "dip" of the graph of every mechanism.

Instead, in this paper, we prove a slightly different figure. We introduce a sufficient condition called Condition M, which is a condition on the distribution of bidders' *valuations* (recall that the auctioneer has confidence in the distribution of bidders' valuations although not in the distribution of bidders' beliefs). We prove that, under Condition M, Figure 2 will be true: there will be a particular assumption about (or distribution of) bidders' beliefs, at which point the graph of *every* (potentially very complicated) Bayesian incentive compatible mechanism must dip below the graph of the best dominant strategy mechanism.

As we will see, Condition M is simply a generalization of what Myerson (1981) calls the "regular case" in his classical paper on optimal auctions. It is a familiar condition in the literature of mechanism design and comfortably assumed in many applications. Hence we consider proving Figure 2 under Condition M as a substantial progress towards the ultimate goal of proving Figure 1.

Section 2 presents the model and formalizes the problem. Section 3 then uses a two-bidder two-valuation example to illustrate our proof of Figure 2. In any two-bidder two-valuation example, Condition M will be satisfied as long as there is no unambiguous strong bidder; i.e., bidder 1's low valuation is lower than bidder 2's high valuation, and vice versa. Our main result will be presented and proved in Section 4. In Section 5, we shall make some remarks on the common prior assumption.

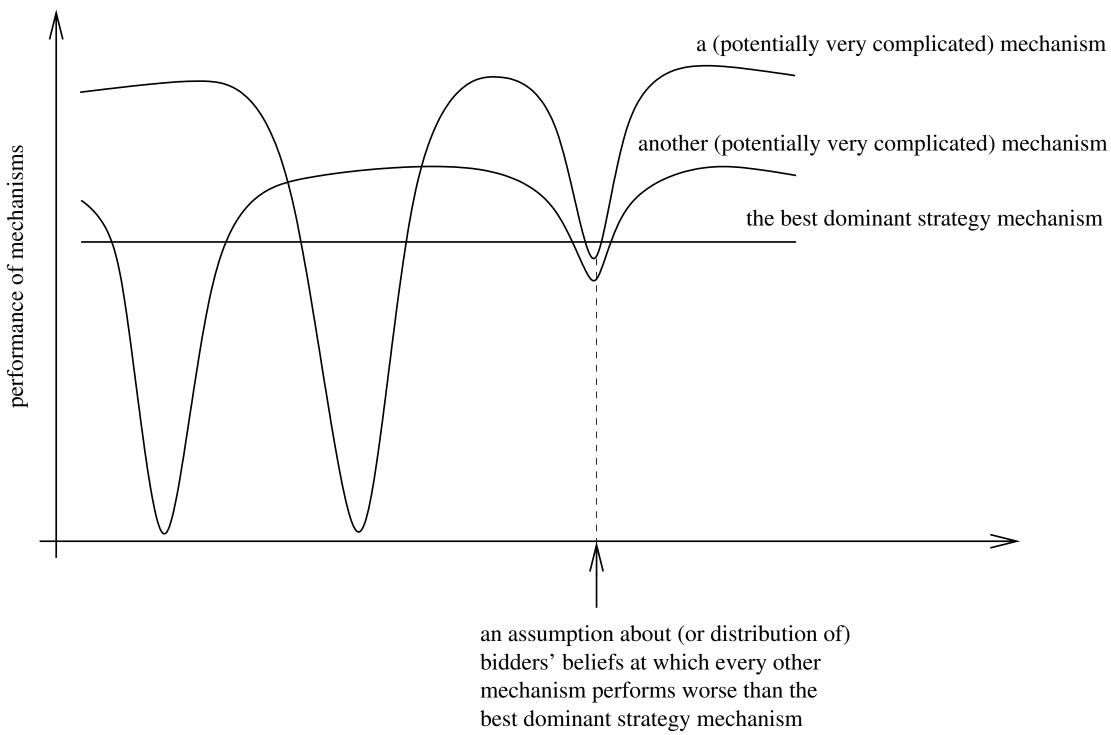

Figure 2: there is a particular point at which the graph of every mechanism dips below the graph of the best dominant strategy mechanism.

This paper is not the first one trying to offer a foundation for dominant strategy mechanisms. Bergemann and Morris (2002) offers an alternative foundation for ex post incentive compatible mechanisms, which in private-value settings are equivalent to dominant strategy mechanisms. It will be easier to compare this paper with Bergemann and Morris (2002) after we present the model. Therefore we shall defer the discussion to Section 7.

Nor would this paper and Bergemann and Morris (2002) be the last ones trying to offer foundations for dominant strategy mechanism. We consider different foundations not as competitors, but instead as complements in building our collective confidence in a particular approach to implement the Wilson Doctrine. Indeed, in Section 6, we investigate yet another possible foundation called the Bayesian foundation. The Bayesian foundation can be loosely explained with the following story. Imagine the auctioneer as a Bayesian decision maker. When she needs to choose a mechanism under uncertainty of bidders' beliefs, she forms a subjective belief about bidders' beliefs, and compares different mechanisms by calculating the expected performance with respect to that subjective belief. When we as outside observers observe that this auctioneer chooses a particular mechanism, we can ask whether or not such a choice is consistent with Bayesian rationality; i.e., whether or not such a choice is optimal with respect to *some* subjective beliefs. If the answer is affirmative, then we say that such a choice is rationalizable. We can say that dominant strategy mechanisms are rationalizable if they are optimal with respect to *some* subjective beliefs.

The difference between the Bayesian and the maxmin foundations is hence whether the auctioneer is a Bayesian or a maxmin decision maker. Given the predominant role of Bayesian rationality in the literature of mechanism design, it seems even more natural to pursue the Bayesian foundation. However, in Section 6, we show that such a foundation is impossible. Indeed, it is not difficult to find examples where dominant strategy mechanisms cannot be rationalized by *any* subjective beliefs.

Section 7 then concludes the paper.

2 Preliminaries

2.1 Types

A single unit of an indivisible object is up for sale. There are N risk-neutral bidders with privately known valuations competing for the object. Each bidder has M possible valuations and for notational simplicity, we suppose that the set V_i of possible valuations is the same for each bidder i and that $V_i = \{v^1, v^2, \ldots, v^M\}$ where $v^m - v^{m-1} = \Delta$ for each m.[1] A bidder i with valuation v_i receives expected utility $p_i v_i - t_i$ if p_i is the probability with which he will be awarded the object and if his expected monetary payment is t_i.

A typical element of $V := \times_i V_i$ is v, and a typical element of $V_{-i} := \times_{j \neq i} V_j$ is v_{-i}.

To characterize the (equilibrium) behavior of the bidders who compete in some given auction mechanism, it is not enough to specify the bidders' possible payoff-relevant types or even the probability distribution from which they are drawn. In addition, we must also specify their beliefs about the valuations of their opponents (called the *first-order* beliefs), their beliefs about one another's' first-order beliefs (*second-order* beliefs), etc.

We wish to consider a formulation of the optimal auction problem which avoids implicit assumptions on higher-order beliefs. The way to do this is to first consider the *universal* belief space in which for every conceivable (coherent) hierarchy of higher-order beliefs there is a representative "belief type." This prevents the modeler from implicitly building in any assumptions about the connections between beliefs among bidders and across orders. Then a "type" consists of a payoff-relevant type together with a belief type. The universal type space is the set of all such types. Finally, we model any assumption about bidders' types (including any possible common knowledge assumption) as a probability measure over the universal type space.

Specifically, we construct the universal belief space from the basic payoff-relevant data as follows (the construction is standard, see Mertens and Zamir (1985) for the details and Brandenburger and Dekel (1993) for an alternative derivation). To begin with, whenever X is a metric space, we treat X as a measurable space with the Borel σ-algebra and let ΔX

[1] These notational conventions simplify the statements of results and notation, but are entirely innocuous. Assumptions of asymmetry in the bidders' valuation sets, or differing gaps between valuations, can be built into the distribution of valuations by assuming that certain valuation profiles have zero probability.

be the space of all Borel probability measures on X endowed with the weak topology.

The set of possible first-order beliefs for bidder i is

$$\mathcal{T}_i^1 := \Delta V_{-i},$$

and the set of all possible kth-order beliefs is

$$\mathcal{T}_i^k := \Delta(V_{-i} \times \mathcal{T}_{-i}^{k-1}).$$

Because the set ΔX is compact metric whenever X is, by induction each $\mathcal{T}_i^k$ is a compact metric space. The projections $\phi_i^k : \mathcal{T}_i^k \to \mathcal{T}_i^{k-1}$, defined inductively by $\phi_i^2(\tau_i^2)(v_{-i}) = \tau_i^2(\{v_{-i}\} \times \mathcal{T}_{-i}^1)$, and for each measurable subset $\{v_{-i}\} \times B \subset V_{-i} \times \mathcal{T}_{-i}^{k-2}$,

$$\phi_i^k(\tau_i^k)(\{v_{-i}\} \times B) = \tau_i^k(\{v_{-i}\} \times \left[\phi_{-i}^{k-1}\right]^{-1}(B)),$$

demonstrate that each kth-order belief for i implicitly defines beliefs at lower orders as well.

A universal belief type for bidder i is a sequence (or *hierarchy*) $\tau_i = (\tau_i^1, \tau_i^2, \ldots)$ satisfying $\tau_i^k \in \mathcal{T}_i^k$ and the *coherency* condition that $\phi_i^k(\tau_i^k) = \tau_i^{k-1}$. The universal belief space for bidder i is then the set $\mathcal{T}_i^* \subset \prod_{k=1}^{\infty} \mathcal{T}_i^k$ of all such coherent hierarchies. This product space endowed with the product topology is compact. Since the set of coherent hierarchies is closed, the universal belief space is compact. By Mertens and Zamir (1985) and Brandenburger and Dekel (1993), there is a homeomorphism between $\mathcal{T}_i^*$ and $\Delta(V_{-i} \times \mathcal{T}_{-i}^*)$ and thus the latter is compact. Let $g_i : \mathcal{T}_i^* \to \Delta(V_{-i} \times \mathcal{T}_{-i}^*)$ be such a mapping.

A *type* is a pair $\omega_i = (v_i, \tau_i)$. Let $f_i(\omega_i) = v_i$ be the projection from bidder i's type to bidder i's valuation. A type space is a set $\Omega = \prod_{i=1}^N \Omega_i$, where $\Omega_i \subset V_i \times \mathcal{T}_i^*$. In this paper, we will mainly deal with two varieties of type spaces.

The *universal type space* Ω^* is the type space where each $\Omega_i^* = V_i \times \mathcal{T}_i^*$. Notice that every type space is a subset of the universal type space. Let $\mathcal{T}^* = \prod_{i=1}^N \mathcal{T}_i^*$. For any $v \in V$, we shall write $\Omega^*(v)$ for the open subset $\{v\} \times \mathcal{T}^* \subset \Omega^*$.

Another kind of type space, used almost without exception in the literature of mechanism design, is the *naive type space* Ω^ν generated from some distribution ν over the set of payoff-relevant types V. Specifically, this means that bidder i's first-order belief is a function of his valuation v_i and is given by the conditional distribution $\tau_i^1(v_i)(\cdot) = \nu(\cdot|v_i)$. Furthermore, since for each $j \neq i$, bidder j's first-order belief is $\tau_j^1(v_j)(\cdot) = \nu(\cdot|v_j)$, bidder i's second-order beliefs can be computed from ν as well. In particular, bidder i believes that with probability

$$\tau_i^2(v_i)(\gamma) := \nu(\left[\tau_{-i}^1\right]^{-1}(\gamma)|v_i),$$

bidders $-i$ have first-order beliefs γ. Similarly, all higher-order beliefs can be inductively derived from ν.

2.2 Mechanisms

We consider direct revelation mechanisms.[2] A direct revelation auction mechanism for type space Ω is a game form in which the bidders simultaneously announce their types from the corresponding set Ω_i, and the object is allocated and monetary transfers enforced as a function of their announcements. Formally, an auction mechanism $\Gamma = (p, t)$ is defined by two functions, $p : \Omega \to [0, 1]^N$ and $t : \Omega \to \mathbf{R}^N$. The allocation rule p specifies the probabilities $p_i(\omega)$ with which each bidder i will receive the object. The allocation rule is restricted to be feasible: $\sum_{i=1}^{N} p_i(\omega) \leq 1$. The transfer rule t defines payments $t_i(\omega)$ made from bidder i to the auctioneer. Denote by $\bar{t}(\omega)$ the sum $\sum_{i=1}^{N} t_i(\omega)$.

We want to show that a cautious auctioneer will choose a simple, dominant strategy incentive compatible auction over more complicated, Bayesian incentive compatible auctions that ask bidders to report their universal types. So we shall define these two notions of incentive compatibility now. Note that, to avoid being too mouthful, we slightly abuse commonly used terminology below.

Definition 1 *An auction mechanism Γ is dominant strategy incentive compatible with respect to the naive type space Ω^ν (or simply dsIC) if for each bidder i and type profile $\omega \in \Omega^\nu$,*

$$\begin{aligned} &p_i(\omega)v_i - t_i(\omega) \geq 0, \qquad \textit{and} \\ &p_i(\omega)v_i - t_i(\omega) \geq p_i(\hat{\omega}_i, \omega_{-i})v_i - t_i(\hat{\omega}_i, \omega_{-i}), \end{aligned}$$

for any alternative type $\hat{\omega}_i \in \Omega_i^\nu$.

Since $|\Omega_i^\nu| = |V_i|$, and since the incentive compatibility constraints for dsIC depend only on valuations, an auction mechanism is dsIC with respect to a naive type space Ω^ν if and only if it is dsIC with respect to any other naive type space $\Omega^{\nu'}$. So we can always discuss whether an auction mechanism is dsIC with respect to the naive type space without referring to the specific distribution ν from which the naive type space is generated.

Definition 2 *An auction mechanism Γ is Bayesian incentive compatible with respect to the universal type space Ω^* (or simply BIC) if for each bidder i and type $\omega_i \in \Omega_i^*$,*

$$\begin{aligned} &\int_{\Omega_{-i}^*} [p_i(\omega)v_i - t_i(\omega)] g_i(\tau_i)(d\omega_{-i}) \geq 0, \qquad \textit{and} \\ &\int_{\Omega_{-i}^*} [p_i(\omega)v_i - t_i(\omega)]\, g_i(\tau_i)(d\omega_{-i}) \geq \int_{\Omega_{-i}^*} [p_i(\hat{\omega}_i, \omega_{-i})v_i - t_i(\hat{\omega}_i, \omega_{-i})]\, g_i(\tau_i)(d\omega_{-i}), \end{aligned}$$

for any alternative type $\hat{\omega}_i \in \Omega_i^*$.

[2] The revelation principle can be shown by standard arguments to hold for all type spaces and all definitions of incentive compatibility considered here. It is thus without loss of generality to restrict attention to direct revelation mechanisms.

Note that any auction mechanism Γ that is dominant strategy incentive compatible with respect to the naive type space (i.e., dsIC) can be extended naturally into an auction mechanism that is Bayesian incentive compatible with respect to the universal type space (i.e., BIC) in a straightforward manner. We shall abuse notation and use Γ to denote this natural extension as well.

2.3 Digression: Why Do We Treat IR and IC "Equally?"

We want to make an important remark on why we treat individual rationality (IR) and incentive compatibility (IC) "equally." In particular, we incorporate the ex post IR constraint into our definition of dominant strategy incentive compatibility, and incorporate the interim IR constraint into our definition of Bayesian incentive compatibility. Is it appropriate to ignore other possible combinations?

In any study on foundations, such as this paper, it is important to reflect on the origins of these IR and IC constraints; i.e., what are the primitive constraints on available mechanisms that are translated into these IC and IR constraints?

There are two kinds of primitive constraints. The first primitive constraint is concerned of the following institutional question: does the mechanism designer have or have not the power to force agents to play her mechanism? If the answer is no, then there will be a constraint that says every mechanism must include an "opt out" option for every agent. If the answer is yes, which may be the case if the mechanism designer is a mafia boss, then there will be no such constraint.

The second primitive constraint is concerned of the solution concept used by the mechanism designer. Every mechanism induces an incomplete information game, and the mechanism designer needs some solution concept to predict the performance of the mechanism and choose one mechanism out of many. Once a solution concept $\mathcal{E}$ is specified, there will be a constraint that says every mechanism must possess an $\mathcal{E}$-equilibrium.

The IR constraint comes from both the first and the second primitive constraints: if the mechanism designer is not a mafia boss, and if she uses dominant strategy equilibrium (respectively Bayesian Nash equilibrium) as her solution concept, then the corresponding IR constraint will be ex post IR (respectively interim IR).

Note that the IC constraint also come from the same second primitive constraint. It means the IR and IC constraints must be treated "equally." Had we used a combination of, say, interim IR and dominant strategy incentive compatibility, we would have difficulty in backing out the second primitive constraint.

There is also a third kind of IR constraint, namely ex ante IR, which we do not consider in this paper. Whether ex ante IR is relevant crucially depends on whether or not there is indeed an ex ante stage. The existence of such a stage is a subject of debate (see, for example, Gul (1998)). For some authors, we included, the existence of an intermediate stage (somewhere in between the ex ante and the interim stages) at which agents receive "signals" or "information" about their own beliefs is inconceivable.

2.4 The Auctioneer as a Maxmin Decision Maker

We model any *assumption* (including any common knowledge assumption) about bidders' belief as a distribution over bidder's universal type space. Specifically, let μ be a distribution over Ω^*. For any BIC auction Γ, the performance of Γ under assumption μ, or the μ-expected revenue, is defined as $R_\mu(\Gamma) = \int_{\Omega^*} \bar{t}\, \mu(d\omega)$.

We take the distribution ν over V as given. This represents the auctioneer's estimate of the bidders' valuations. An assumption that is consistent with this estimate is a distribution μ on the universal type space Ω^* whose marginal on V is ν. Let $\mathcal{M}(\nu)$ denote the compact subset of such assumptions. Observe that there is a unique element ν^* in this subset that concentrates on the naive type space Ω^ν generated by ν. This represents the (common knowledge) assumption in the traditional literature that Wilson (1987) refers to. Unlike the standard formulation of the optimal auction design problem, we do not assume that the auctioneer has confidence in this particular assumption ν^*. Rather the auctioneer considers other assumptions within the set $\mathcal{M}(\nu)$ as possible as well.

A cautious auctioneer who chooses an auction that maximizes the worst-case performance is hence solving the *maxmin* problem of

$$\sup_{\Gamma \text{ is BIC}} \inf_{\mu \in \mathcal{M}(\nu)} R_\mu(\Gamma). \tag{1}$$

Note that if an auction Γ is dominant strategy incentive compatible with respect to the naive type space (i.e., dsIC), then for any assumption $\mu \in \mathcal{M}(\nu)$, the μ-expected revenue of Γ—or, more precisely, Γ's natural extension into the universal type space—depends only on the distribution ν. Hence we can write $R_\mu(\Gamma)$ as $R_\nu(\Gamma)$ without ambiguity.

Definition 3 *Given any distribution ν over V, the optimal dsIC revenue is defined as*

$$\Pi^D(\nu) := \sup_{\Gamma \textit{ is dsIC}} R_\nu(\Gamma).$$

The maxmin foundation of dominant strategy mechanisms refers to the following equation:

$$\Pi^D(\nu) = \sup_{\Gamma \text{ is BIC}} \inf_{\mu \in \mathcal{M}(\nu)} R_\mu(\Gamma), \tag{2}$$

for every distribution ν over valuations.

In this paper, instead of proving that equation (2) holds for every ν, we shall prove that it hold for every ν satisfying a sufficient condition called Condition M (to be defined in Section 4). Specifically, we shall prove that, whenever ν satisfies Condition M, there will exist an assumption $\mu^* \in \mathcal{M}(\nu)$ consistent with ν under which we will have

$$\Pi^D(\nu) = \sup_{\Gamma \text{ is BIC}} R_{\mu^*}(\Gamma), \tag{3}$$

which implies

$$\Pi^D(\nu) = \sup_{\Gamma \text{ is BIC}} R_{\mu^*}(\Gamma) \geq \inf_{\mu \in \mathcal{M}(\nu)} \sup_{\Gamma \text{ is BIC}} R_\mu(\Gamma) \geq \sup_{\Gamma \text{ is BIC}} \inf_{\mu \in \mathcal{M}(\nu)} R_\mu(\Gamma) \geq \inf_{\mu \in \mathcal{M}(\nu)} \Pi^D(\nu) = \Pi^D(\nu),$$

or simply

$$\Pi^D(\nu) = \sup_{\Gamma \text{ is BIC}} \inf_{\mu \in \mathcal{M}(\nu)} R_\mu(\Gamma),$$

which delivers the maxmin foundation as promised.

3 An Illustrative Example

In this section, we shall use a simple example to illustrate our main result as well as the strategy of proof.

Consider an auction example with two bidders, and each bidders have two possible valuations. Bidders' valuations are correlated according to the distribution ν depicted in Figure 3.

	$v_1 = 4$	$v_1 = 9$
$v_2 = 11$	3/10	1/10
$v_2 = 5$	3/10	3/10

Figure 3: The distribution ν of bidders' valuations.

The optimal dsIC auction is depicted in Figure 4. In Figure 4, "$\alpha = i$" is the shorthand for "allocating the object to bidder i" (i.e., $p_i = 1$ and $p_{-i} = 0$), and "$\alpha = 0$" means no sale.

	$v_1 = 4$	$v_1 = 9$
$v_2 = 11$	$\alpha = 2, t_1 = 0, t_2 = 11$	$\alpha = 2, t_1 = 0, t_2 = 11$
$v_2 = 5$	$\alpha = 0, t_1 = 0, t_2 = 0$	$\alpha = 1, t_1 = 9, t_2 = 0$

Figure 4: The optimal dsIC auction Γ.

In any two-bidder two-valuation example, Condition M (to be formally defined in the next section) will be satisfied as long as there is no unambiguous strong bidder; i.e., bidder 1's low valuation is lower than bidder 2's high valuation, and vice versa. Hence, according to Theorem 1 (to be proved in the next section), there exists an assumption μ^* consistent with the distribution ν such that equation (3) holds.

We construct one such assumption μ^* below, but shall keep our exposition informal. Let a_i (b_i) denote the first-order belief of a high-valuation (low-valuation) type of bidder i that bidder $-i$ has high valuation.

Consider an assumption μ^* which has a 4-point support: for every bidder i, every possible valuation is associated with only one possible belief type. The marginal distribution of μ^* over bidders' valuations and first-order beliefs is as depicted in Figure 5.

	$b_1 = 2/5$	$a_1 = 1/4$
$a_2 = 1/4$	3/10	1/10
$b_2 = 2/5$	3/10	3/10

Figure 5: The auctioneer's belief μ.

The bidders' higher-order beliefs are derived from Figure 5 by induction. For example, for a low-valuation type of bidder 1, his second-order belief assigns probability 2/5 (3/5) to bidder 2 having high (low) valuation and holding first-order belief $a_2 = 1/4$ ($b_2 = 2/5$), and a high-valuation (low-valuation) type of bidder 2 has a third-order belief that assigns probability 3/4 (3/5) to bidder 1 having low valuation and having such a second-order belief, and so on.

It is obvious that this assumption μ^* is consistent with the distribution ν.

Under this assumption μ^*, there are at least two possible ways to improve upon the optimal dominant strategy auction Γ in Figure 4. First, according to the assumption μ^*, conditional on bidder 1 having low valuation, the conditional probability that bidder 2 has high valuation is 1/2. This is different from the first order belief of the low-valuation type of bidder 2, which is $b_1 = 2/5$. So one possible way to improve upon Γ is to betting against the low-valuation type of bidder 1 on bidder 2's types. Second, since high- and low-valuation types of bidder 1 hold different beliefs, another possible way to improve upon Γ is to separate these two types by introducing Crémer-McLean-kind of lotteries and relaxing incentive compatibility constraints. We shall see that either of these two ways fail to improve upon Γ.

First, consider introducing any bet (x, y) on bidder 2's type, where x and y are the amount bidder 1 pays the auctioneer in the events bidder 2 has low and high valuations respectively. If the bet is acceptable to both the auctioneer and the low-valuation type of bidder 1, we must have

$$\begin{aligned}(1/2)x + (1/2)y &\geq 0, \quad \text{and} \\ (3/5)(-x) + (2/5)(-y) &\geq 0,\end{aligned}$$

with at least one inequality strict unless $x = y = 0$. But then the high-valuation type of bidder 1 would find the bet acceptable as well, as

$$(3/4)(-x) + (1/4)(-y) = (5/2)[(3/5)(-x) + (2/5)(-y)] + (3/2)[(1/2)x + (1/2)y],$$

which is strictly bigger then the zero rent for the high-valuation type of bidder 1 under the auction Γ. With both high- and low-valuation types of bidder 1 accepting such a bet, such a bet turns sour for the auctioneer, as

$$(3/5)(-x) + (2/5)(-y) \leq 0,$$

and this explains why introducing the first kind of bets does not help.

Second, consider introducing any Crémer-McLean-kind of lottery to separate the high- and low-valuation types of bidder 1 and relax the downward incentive compatibility constraint. Once again, let (x, y) be such a bet on bidder 2's type. Suppose the bet is successful in the sense that the auctioneer can now sell to the low-valuation type of bidder 1 without the need to leave extra rent for the high-valuation type of bidder 1 (as she needed to before the introduction of such a bet that relaxes the downward incentive compatibility constraint), then we must have

$$\begin{aligned} &(3/5)(4-x) + (2/5)(-y) \geq 0, \quad \text{and} \\ &(3/4)(9-x) + (1/4)(-y) \leq 0, \end{aligned}$$

where the first (second) inequality follows from the individual rationality (incentive compatibility) constraint of the low-valuation (high-valuation) type of bidder 1. However, these together imply that any bet like this is too good to be profitable for the auctioneer, as

$$(1/2)x + (1/2)y = (2/3)[(3/4)(-x) + (1/4)(-y)] - (5/3)[(3/5)(-x) + (2/5)(-y)] \leq -1,$$

and this explains why introducing the second kind of bets does not help either.

In principle, there may still be other possible ways to improve upon the optimal dsIC auction Γ. But actually there are no more (this requires a proof, which will be the content of Theorem 1). Hence, under the assumption μ^* equation (3) holds, which in turn implies that equation (2) holds as well.

4 The Main Result

In this section, we shall first review the problem optimal dsIC auction design. We use a version of a standard argument to show that the dominant strategy incentive constraints can be replaced by a monotonicity constraint on the allocation rule. We then formally define Condition M, which in effect says the monotonicity constraint is not binding in the optimal dsIC auction design problem. We will relate Condition M to the more familiar condition called "the regular case" in the optimal auctions literature after Myerson (1981), and show in Proposition 1 sufficient conditions (analogous to the more familar "monotone hazard rate" condition) under which Condition M holds.

4.1 Review of Optimal dsIC Auctions

We can formulate the optimal dsIC auction design problem as follows:

$$\max_{p(\cdot),t(\cdot)} \quad \sum_{v_i \in V} \nu(v) \sum_{i=1}^{N} t_i(v) \tag{4}$$

$$\text{subject to:} \quad \forall\, i = 1,\dots,N,\ \forall\, m,l = 1,\dots,M,\ \forall\, v_{-i} \in V_{-i},$$

$$p_i(v^m, v_{-i})v^m - t_i(v^m, v_{-i}) \geq 0, \qquad \langle DIR_i^m \rangle$$

$$p_i(v^m, v_{-i})v^m - t_i(v^m, v_{-i}) \geq p_i(v^l, v_{-i})v^m - t_i(v^l, v_{-i}). \qquad \langle DIC_i^{m \to l} \rangle$$

By some standard manipulations, we shall eliminate some constraints and rewrite the problem in a form that will facilitate comparison with the optimal BIC auction. The following result is standard.

Lemma 1 *Say that an allocation rule p is dsIC if there exists a transfer rule t such that the auction mechanism (p,t) satisfies the constraints in (4). A necessary and sufficient condition for p to be dsIC is the following monotonicity condition:*

$$p_i(v^m, v_{-i}) \geq p_i(v^{m-1}, v_{-i}), \qquad \forall\, m = 2,\dots,M, \qquad \forall\, v_{-i} \in V_{-i}. \qquad \langle M_i \rangle$$

It follows again from standard arguments that in an optimal dsIC auction, the constraints $\langle DIR_i^1 \rangle$ and $\langle DIC_i^{m \to m-1} \rangle$ are binding and (given that p is monotonic) all other constraints can be ignored. Combining the resulting two equalities, we see that when the other bidders report valuation profile v_{-i}, bidder i's net utility ("rent") will be

$$U_i(v^1, v_{-i}) = 0$$

for type v^1 and

$$U_i(v^m, v_{-i}) = p_i(v^{m-1}, v_{-i})(v^m - v^{m-1}) + U_i(v^{m-1}, v_{-i}) = \Delta \sum_{m'=1}^{m-1} p_i(v^{m'}, v_{-i})$$

for type v^m, $m > 1$. By definition, the total transfer received by the auctioneer is the total surplus generated by any sale of the object less the rent received by the bidders. Thus, an equivalent formulation of the problem is to choose a dsIC (i.e., monotonic) allocation rule to maximize the expected value of this difference:

$$\max_{p(\cdot)} \quad \sum_{i=1}^{N}\sum_{m=1}^{M}\sum_{v_{-i}\in V_{-i}} \nu(v^m, v_{-i}) \left[p_i(v^m, v_{-i})v^m - \Delta \sum_{m'=1}^{m-1} p_i(v^{m'}, v_{-i})\right] \qquad (5)$$

$$\text{subject to } \langle M_i \rangle,\ i = 1, \ldots, N.$$

In accordance with Lemma 1, the monotonicity constraint appears as an equivalent expression for dsIC. This constraint may or may not bind at the solution. The literature on optimal auctions has developed various techniques for incorporating the constraint when it does bind. Loosely speaking, one first ignore the monotonicity constraint and solves the relaxed problem by pointwise optimization. If the resulting allocation rule is already monotonic, then it is also a solution to the constrained problem. If the resulting allocation rule is not monotonic, some extra "ironing" work will be needed. In many applications, the "ironing" work is typically avoided by imposing certain conditions (such as "monotone hazard rate") on the distribution ν over valuations. The case where "ironing" is not needed is generally referred to as "the regular case."

Condition M *We say that the distribution ν over bidders' valuations satisfies Condition M if the constraints $\langle M_i \rangle$ are not binding in problem* (5).

Note for future use that the set of ν satisfying Condition M is closed. Indeed, if ν satisfies Condition M, the solution to problem (5) can be solved by pointwise optimization, which in turn will be continuous in ν. That the set of ν is closed then follows from the fact that the set of monotonic allocation rules is closed.

Before we prove that equation (2) holds for all ν satisfying Condition M, let us point out conditions that are in turn sufficient for Condition M. We should emphasize that, in the traditional literature, bidders' valuations are typically assumed to be independently distributed. With independent types, sufficient conditions for Condition M are well-known. In our setting, bidders' valuations may be correlated, and this makes Condition M slightly more complicated. Let $F_i(v^m, v_{-i}) = \sum_{m'=m+1}^{M} \nu(v^{m'}, v_{-i})$ denote the de-cumulative distribution function for bidder i's valuation, given the valuations v_{-i} of the other bidders. The inverse *hazard rate* for bidder i is $h_i(v) = F_i(v)/\nu(v)$.

Proposition 1 *A distribution ν over bidders' valuations will satisfy Condition M if, for each i, j, and v_{-i}, the difference $h_i(v_i, v_{-i}) - h_j(v_i, v_{-i})$ is decreasing in v_i.*

Note that in the case of independent ν, our condition reduces to the usual condition of monotone hazard rate. For more general ν, our condition is neither weaker nor stronger than monotone hazard rate. It is always satisfied when bidders' valuations are sufficiently positively correlated. Proposition 1 is proved in the appendix. We now turn to our main result.

4.2 The Possibility of Maxmin Foundation

Theorem 1 *Equation (2) holds for all ν satisfying Condition M.*

Proof: Let ν be given. Write ν_i^m for the marginal probability of valuation $v_i = v^m$, and write $F_i(m) = \sum_{m'=m}^{M} \nu_i^{m'}$ for the associated de-cumulative distribution function. Let $\sigma_i^m = \nu(\cdot|v^m)$ be the conditional distribution over the valuations of bidders $j \neq i$ conditional on bidder i having valuation v^m. We first analyze the case in which ν is *regular*; i.e., $\nu_i^m > 0$ for each m (so that these conditional distributions are well-defined) and the collection of vectors $\{\sigma_i^m\}_{m=1}^M$ is linearly independent.

Our proof is constructive. Construct an assumption μ^* which concentrates on M possible types for each bidder. Let $\Omega = \times_i \Omega_i$ be the support of μ^*, with $\Omega_i = \{\omega_i^m\}_{m=1}^M = \{(v^m, \tau_i^m)\}_{m=1}^M$ representing the set of possible types of bidder i under assumption μ. The beliefs τ_i^m of these types will be specified next. For each $\omega_j \in \Omega_j$, let $f_j(\omega_j)$ be the valuation of ω_j. Note that f_j is a bijection for all j. For any belief τ over V_{-i}, define a corresponding belief $\pi_i(\tau)$ over Ω_{-i} in the straightforward way: $\pi_i(\tau)(\omega_{-i}) = \tau((f_j(\omega_j))_{j\neq i})$. In what follows, we shall occasionally use the notation τ interchangeably for $\pi_i(\tau)$, and the context will prevent any confusion.

We construct the bidders' beliefs as follows:

$$\forall i, \forall m, \quad \tau_i^m = \frac{1}{F_i(m)} \sum_{m'=m}^{M} \nu_i^{m'} \sigma_i^{m'}.$$

Thus, conditional on having valuation v^m, bidder i's belief over opponents' valuations (and hence types) is a conditional expectation with respect to ν; in particular, it is the average ν-probability conditional on i having valuation *at least* v^m.[3] Note that the collection $\{\tau_i^m\}_{m=1}^M$ is linearly independent. The following equivalent recursive definition of τ_i^m is useful:

$$\begin{aligned} \tau_i^M &= \sigma_i^M, \\ \tau_i^m &= \frac{1}{F_i(m)} \left(\nu_i^m \sigma_i^m + F_i(m+1)\tau_i^{m+1}\right), \quad \forall m < M. \end{aligned} \tag{6}$$

Finally, we specify the assumption μ^* about types: $\mu^* = \pi(\nu)$; i.e., $\mu^*(\omega) = \nu((f_i(\omega_i)_{i=1}^N)$. Obviously $\mu^* \in \mathcal{M}(\nu)$. Under this assumption μ^*, the optimal BIC auction design problem

[3] Thus, each bidder type has beliefs which are a distortion of ν, except for the highest valuation type, where there is "no distortion at the top."

is as follows:

$$\max_{p(\cdot),t(\cdot)} \quad \sum_{i=1}^{N} \sum_{\omega\in\Omega} \mu^*(\omega) t_i(\omega) \tag{7}$$

$$\text{subject to:} \quad \forall\ i = 1,\dots,N,\ \forall\ m = 1,\dots,M,\ \forall\ l = 1,\dots,M,$$

$$\tau_i^m \cdot (p_i^m v^m - t_i^m) \geq 0, \qquad \langle IR_i^m \rangle$$

$$\tau_i^m \cdot (p_i^m v^m - t_i^m) \geq \tau_i^m \cdot \left(p_i^l v^m - t_i^l\right). \qquad \langle IC_i^{m\to l} \rangle$$

We have used the shorthand notation p_i^m and t_i^m to refer to the vectors $p_i(\omega_i^m, \cdot)$ and $t_i(\omega_i^m, \cdot)$ respectively in $\mathbf{R}^{M^{N-1}}$, and the inner product notation such as $\tau_i^m \cdot p_i^m$ for the expectations of these vectors with respect to the belief τ_i^m.

Say that an allocation rule p is BIC if there exists a transfer rule t such that the auction mechanism (p, t) satisfies the constraints in (7). Because the beliefs of the types of each bidder are linearly independent, every allocation rule is BIC. Indeed, by exploiting the differences in beliefs, the incentive compatibility and individual rationality constraints can be satisfied by building into the transfer rule lotteries which have positive expected value to the intended type and arbitrarily large negative expected values to the other types. This kind of construction is due to Crémer and McLean (1988), and we shall omit the details.

While the above argument shows that any allocation rule is implementable by some appropriate choice of transfer rule, we can further sharpen the conclusion and argue that certain constraints in (7) can be manipulated or even ignored without cost to the auctioneer. To begin with, each "upward" incentive constraint (i.e., $\langle IC_i^{m\to l}\rangle$ for $m < l$) can be ignored. Indeed, because bidder i's beliefs are linearly independent, there exists a lottery $\lambda \in \mathbf{R}^{M^{N-1}}$ such that $\tau_i^m \cdot \lambda = 0$ for all $m \geq l$ and $\tau_i^m \cdot \lambda < 0$ for all $m < l$. Since by (6) σ_i^l is a linear combination of τ_i^l and τ_i^{l+1}, we also have $\sigma_i^l \cdot \lambda = 0$. By adding (some sufficiently large scale of) λ to t_i^l, each $\langle IC_i^{m\to l}\rangle$ for $m < l$ can be relaxed. No other constraints are affected and the resulting change in the auctioneer's revenue is $\sigma_i^l \cdot \lambda = 0$.

We next show that for any auction mechanism (p, t) that satisfies the remaining constraints, there exists an auction mechanism (p', t') which satisfies the constraints $\langle IR_i^m\rangle$, for $m = 1, \dots, M$, and $\langle IC_i^{m\to m-1}\rangle$, for $m = 2, \dots, M$, with equality, and achieves at least as high an μ^*-expected revenue as (p, t) does.

To prove this, fix any auction mechanism (p, t) that satisfies the remaining constraints. Suppose $\langle IC_i^{m\to m-1}\rangle$ holds with strict inequality. Let $\boldsymbol{\tau}$ denote the matrix whose M rows are the vectors $\{\tau_i^m\}_{m=1}^M$, and let $(\boldsymbol{\tau}^{-m}, \sigma_i^{m-1})$ be the matrix obtained by replacing the mth row of $\boldsymbol{\tau}$ with the vector σ_i^{m-1}. Note that the matrix $(\boldsymbol{\tau}^{-m}, \sigma_i^{m-1})$ has rank M. We can thus solve the following equation for λ:

$$(\boldsymbol{\tau}^{-m}, \sigma_i^{m-1}) \cdot \lambda = \mathbf{e}^m,$$

where $\mathbf{e}^m$ denotes the mth elementary basis vector in $\mathbf{R}^M$. Note that because $\tau_i^{m-1} \cdot \lambda = 0 < \sigma_i^{m-1} \cdot \lambda$, and because τ_i^{m-1} is a convex combination of σ_i^{m-1} and τ_i^m according to (6), we have $\tau_i^m \cdot \lambda < 0$.

We will add the vector $\varepsilon\lambda$ to t_i^{m-1} for some scalar $\varepsilon > 0$. Because $\tau_i^{m'} \cdot \lambda = 0$ for $m' \neq m$, no constraints for types $\omega_i^{m'}$ are affected. As for type ω_i^m, the constraint $\langle IR_i^m \rangle$ is unaffected. The only incentive constraint of type ω_i^m that is affected is $\langle IC_i^{m\to m-1} \rangle$, and this constraint was slack by assumption. Let $S_i^m > 0$ be the slack in $\langle IC_i^{m\to m-1} \rangle$, and choose $\varepsilon = -S_i^m/(\tau_i^m \cdot \lambda) > 0$. Then, with the resulting transfer rule, $\langle IC_i^{m\to m-1} \rangle$ holds with equality. Finally, because $\varepsilon\sigma_i^{m-1} \cdot \lambda > 0$, the auctioneer profits from this modification.

We next show that each $\langle IR_i^m \rangle$ can be treated as an equality without loss of generality. Define $S_i^m = \tau_i^m \cdot (p_i^m v^m - t_i^m) \geq 0$ to be the slack in $\langle IR_i^m \rangle$. Construct a lottery λ that satisfies

$$\tau_i^m \cdot \lambda = S_i^m, \quad m = 1, \ldots, M.$$

By the full-rank arguments such a lottery λ can be found. We will add λ to each t_i^m. No constraint of the form $\langle IC_i^{m\to l} \rangle$ will be affected, but now each constraint of the form $\langle IR_i^m \rangle$ holds with equality. Finally, we check that the auctioneer profits from this modification. Indeed, the auctioneer nets

$$\begin{aligned}
\sum_{m=1}^{M} \nu_i^m \left(\sigma_i^m \cdot \lambda\right) &= \sum_{m=1}^{M-1} \left(F_i(m)\tau_i^m - F_i(m+1)\tau_i^{m+1}\right) \cdot \lambda + \nu_i^M \tau_i^M \cdot \lambda \\
&= F_i(1)\tau_i^1 \cdot \lambda \\
&= F_i(1) S_i^1 \\
&\geq 0.
\end{aligned}$$

The proof for the regular case is now concluded as follows. Based on the preceding arguments, we consider the modified program in which the constraints $\langle IR_i^m \rangle$ and $\langle IC_i^{m\to m-1} \rangle$ are satisfied with equality. We will use these constraints to substitute out for the transfers in the objective function and reduce the problem to an *unconstrained* optimization with the only choice variable being the allocation rule (recall that any allocation rule is BIC). The resulting objective function will be identical to the objective function (4) for the dsIC problem. Thus the only difference between the two problems is the absence of any monotonicity constraint in the BIC case. It then follows that (i) the modified problem and hence the original problem (7) will have a solution, and (ii) this solution will be the same as the solution to the dsIC problem given Condition M.

We rewrite the objective function in (7) as below, and impose the constraints as equalities:

$$\max_{p(\cdot),t(\cdot)} \quad \sum_{i=1}^{N} \sum_{m=1}^{M} \nu_i^m \sigma_i^m \cdot t_i^m \tag{8}$$

$$\begin{aligned}
&\text{subject to:} \quad \forall\, i = 1, \ldots, N,\ \forall\, m = 1, \ldots, M, \\
&\tau_i^m \cdot (p_i^m v^m - t_i^m) = 0, && \langle \overline{IR}_i^m \rangle \\
&\tau_i^m \cdot (p_i^m v^m - t_i^m) = \tau_i^m \cdot \left(p_i^{m-1} v^m - t_i^{m-1}\right). && \langle \overline{IC}_i^{m\to m-1} \rangle
\end{aligned}$$

We have substituted in the objective function using the definition $\mu^*(\omega) = \nu(f(\omega)) =$

$\nu_i(f_i(\omega_i))\nu(f_{-i}(\omega_{-i})|f_i(\omega_i)) = \nu_i^m \sigma_i^m$ for the appropriate m.

By definition, $\sigma_i^M = \tau_i^M$, so $\langle \overline{IR}_i^M \rangle$ becomes $\sigma_i^M \cdot t_i^M = v^M \sigma_i^M \cdot p_i^M$. Now, for arbitrary $m < M$,

$$\begin{aligned}
\sigma_i^m \cdot t_i^m &= \frac{1}{\nu_i^m}\left[F_i(m)\tau_i^m - F_i(m+1)\tau_i^{m+1}\right] \cdot t_i^m \\
&= \frac{1}{\nu_i^m}\left\{F_i(m)v^m\tau_i^m \cdot p_i^m - F_i(m+1)\left[\tau_i^{m+1} \cdot (p_i^m - p_i^{m+1})v^{m+1} + \tau_i^{m+1} \cdot t_i^{m+1}\right]\right\} \\
&= \frac{1}{\nu_i^m}\left[F_i(m)v^m\tau_i^m \cdot p_i^m - F_i(m+1)v^{m+1}\tau_i^{m+1} \cdot p_i^m\right].
\end{aligned}$$

In the first line we used the recursive definition in (6), in the second line we used $\langle \overline{IR}_i^m \rangle$ and $\langle \overline{IC}_i^{m+1 \to m} \rangle$, and in the third line we used $\langle \overline{IR}_i^{m+1} \rangle$.

Substituting the constraints into the objective function, it becomes:

$$\begin{aligned}
& \sum_{i=1}^N \left\{ v^M \nu_i^M \sigma_i^M \cdot p_i^M + \sum_{m=1}^{M-1}\left[v^m F_i(m)\tau_i^m \cdot p_i^m - v^{m+1}F_i(m+1)\tau_i^{m+1} \cdot p_i^m\right]\right\} \\
= & \sum_{i=1}^N \left\{ v^M \nu_i^M \sigma_i^M \cdot p_i^M + \sum_{m=1}^{M-1}\left[v^m \left(\nu_i^m\sigma_i^m + F_i(m+1)\tau_i^{m+1}\right) \cdot p_i^m - v^{m+1}F_i(m+1)\tau_i^{m+1} \cdot p_i^m\right]\right\} \\
= & \sum_{i=1}^N \left[\sum_{m=1}^M v^m \nu_i^m \sigma_i^m \cdot p_i^m - \sum_{m=2}^M (v^m - v^{m-1})F_i(m)\tau_i^m \cdot p_i^{m-1}\right].
\end{aligned}$$

Applying the definition of τ_i^m, the objective function becomes:

$$\begin{aligned}
& \sum_{i=1}^N \left[\sum_{m=1}^M v^m \nu_i^m \sigma_i^m \cdot p_i^m - \sum_{m=2}^M \Delta\left(\sum_{m'=m}^M \nu_i^{m'}\sigma_i^{m'}\right) \cdot p_i^{m-1}\right] \\
= & \sum_{i=1}^N \left[\sum_{m=1}^M v^m \nu_i^m \sigma_i^m \cdot p_i^m - \Delta \sum_{m=2}^M \sum_{m'=2}^m \nu_i^m \sigma_i^m \cdot p_i^{m'-1}\right] \\
= & \sum_{i=1}^N \sum_{m=1}^M \nu_i^m \sigma_i^m \cdot \left[v^m p_i^m - \Delta \sum_{m'=2}^m p_i^{m'-1}\right] \\
= & \sum_{i=1}^N \sum_{m=1}^M \sum_{v_{-i}\in V_{-i}} \nu_i(v^m, v_{-i}) \cdot \left[v^m p_i(v^m, v_{-i}) - \Delta \sum_{m'=1}^{m-1} p_i(v^{m'}, v_{-i})\right].
\end{aligned}$$

This is identical to the objective function in (5). This establishes equation (3), and hence also equation (2), for any regular ν that satisfies Condition M.

Now consider any arbitrary ν, which needs not be regular, that satisfies Condition M. There exists a sequence ν^n converging to ν such that each ν^n is regular and satisfies Condition

M. For each regular ν^n, construct the type space Ω^n exactly as in the first half of the proof. Let $\tau_i^m(n)$ denote the belief of type ω_i^m of bidder i in the type space Ω_i^n. Passing to a subsequence if necessary, take $\tau_i^m(n) \to \tau_i^m$ for each i and m. Let Ω be the limit type space with beliefs τ_i^m, and let $\mu^* = \pi(\nu)$. Write $\mu^n = \pi(\nu^n)$.

Note that each of these type spaces (Ω^n or Ω) has the same property that there is a one-to-one correspondence between types and valuations for each bidder i. Therefore, for any auction mechanism (p, t) defined over any of these type spaces, we can also think of it as mappings from V to probabilities and transfers. The following notations are hence defined regardless of which of these type spaces the auction mechanism (p, t) is defined over:

$$\mathbf{E}_{n'}\bar{t} := \int_V \bar{t}(v)\nu^{n'}(v),$$
$$\mathbf{E}\bar{t} := \int_V \bar{t}(v)\nu(v).$$

Consider any auction mechanism (p, t) that is Bayesian incentive compatible with respect to type space Ω.[4] Obviously $\mathbf{E}_n\bar{t} \to \mathbf{E}\bar{t}$. We will show that there exists a sequence of auction mechanisms $(p, t(n))$ such that each $(p, t(n))$ is Bayesian incentive compatible with respect to type space Ω^n, and such that $\mathbf{E}_n\bar{t}(n) - \mathbf{E}_n\bar{t} \to 0$.

For each i, m, and n, let

$$S_i^m(n) = \max\{0, \tau_i^m(n) \cdot (t_i^m - p_i^m \cdot v^m)\}$$

be the amount by which the $\langle IR_i^m \rangle$ constraint is violated by the auction mechanism (p, t) for type $\omega_i^m(n)$. Because (p, t) is Bayesian incentive compatible with respect to Ω, $S_i^m(n) \to 0$ for each i and m.

However, (p, t) may not be Bayesian incentive compatible with respect to Ω^n. To modify it into an auction mechanism that is Bayesian incentive compatible with respect to Ω^n, we first add the constant $-S_i^m(n)$ to t_i^m to restore all $\langle IR_i^m \rangle$ constraints. The cost of this to the auctioneer is the μ^n-expected value of $S_i^m(n)$ which is converging to zero. Let $\tilde{t}(n)$ be the transfer rule that results from this first step of modification.

Next, for each i, m, l, and n, let

$$L_i^{m\to l}(n) = \max\{0, \tau_i^l(n) \cdot (p_i^l v^m - \tilde{t}_i^l(n)) - \tau_i^m(n) \cdot (p_i^m v^m - \tilde{t}_i^m(n))\}$$

be the amount by which $\langle IC_i^{m\to l} \rangle$ is violated by the auction mechanism $(p, \tilde{t}(n))$. Note that $L_i^{m\to m}(n) = 0$. Again, because (p, t) is Bayesian incentive compatible with respect to Ω, and because $\tilde{t}(n) \to t$, we have $L_i^{m\to l}(n) \to 0$ for each i, m, and l. For each n, we construct $\lambda_i^l(n)$ to solve the system

$$\tau_i^m(n) \cdot \lambda_i^l(n) = L_i^{m\to l}(n), \qquad \forall\, i, m, l.$$

[4]This is similar to the definition of BIC mechanisms (Definition 2) except that Ω^* is replaced with Ω.

We will add $\lambda_i^l(n)$ to $\tilde{t}_i^l(n)$ to restore each $\langle IC_i^{m\to l}\rangle$ constraint, without affecting $\langle IR_i^m\rangle$ constraints. The resulting auction mechanism $(p, t(n))$ is now Bayesian incentive compatible with respect to Ω^n. Since $\tau_i^m(n)\cdot\lambda_i^l(n)\to 0$ for each m and l, so by (6) we have $\sigma_i^m(n)\cdot\lambda_i^m(n)\to 0$ for each m, and thus,

$$\mathbf{E}_n[\bar{t}(n)-\bar{t}]\to 0 \tag{9}$$

as promised.

Finally, recall that $\Pi^D(\nu)$ denotes the optimal dsIC ν-expected revenue. Because the constrained set in the optimal dsIC auction design problem (5) is compact, the maximum theorem implies

$$\Pi^D(\nu^n)\to\Pi^D(\nu). \tag{10}$$

We have already shown that $\Pi^D(\nu^n)\geq\mathbf{E}_n\bar{t}(n)$ because each ν_n is regular. This together with (9) and (10) delivers

$$\mathbf{E}(\bar{t})=\lim_n\mathbf{E}_n\bar{t}\leq\Pi^D(\nu).$$

Since (p,t) was an arbitrary auction mechanism that is Bayesian incentive compatible with respect to Ω, this establishes equation (3), and hence also equation (2), for any ν that satisfies Condition M. ■

5 Remarks on the Common Prior Assumption

In this section, we want to address a few concerns that may arise due to the observations that the particular assumption μ^* we used in the proof of Theorem 1 is an assumption that is at odd with the common prior assumption (CPA).

Loosely speaking, the CPA says that there is a common probability measure (the common prior) from which each bidder derives his belief by computing the conditional probability of opponents' types conditional on his own "signal" or "information." In our current setting, where any assumption about bidders' types is already modelled as a probability distribution over bidders' types, we can relate any assumption μ to the CPA as follows. For any subset $A\in\Omega_i^*$, we shall write $\mu(A)$ as a short hand for $\mu(A\times\Omega_{-i}^*)$. In other words, we abuse notation and use the same notation for a probability measure as well as its marginal distributions. Recall that $g_i:\mathcal{T}_i^*\to\Delta(V_{-i}\times\mathcal{T}_{-i}^*)$ is the homeomorphism between bidder i's belief types and distributions over his opponents' types.

Definition 4 *We say that an assumption μ is an CPA-assumption if for any measurable subsets $A\subset\Omega_i^*$ and $B\subset\Omega_{-i}^*$,*

$$\int_A g_i(\tau_i)(B)\,\mu(d\omega_i)=\mu(A\times B).$$

It is apparent that the particular assumption μ^* we used in the proof of Theorem 1 is not an CPA-assumption. Can we replace μ^* with some CPA-assumption μ in the proof? The answer is: sometimes, but not always. For some distribution ν over bidders' valuations, especially those that are close to being independent, it is indeed possible to use an CPA-assumption in the proof of Theorem 1. But it is also not difficult to find an example of ν such that no such CPA-assumption can be constructed. We will give one such an example in the Appendix.

Do these observations cut back the appeal of Theorem 1? We believe the answer is: not at all, for two reasons. First, whether the CPA is an appropriate assumption to make is itself a subject of debate. Gul (1998) has explained why the CPA lacks appropriate motivations, and Morris (1995) has also explained why many defenses of the CPA are flawed, and why many interesting economic problems are better modelled without the CPA.

Second, recent studies on the CPA has uncovered the close relation between the CPA and common knowledge assumptions (see, for example, Lipman (2003)). In any study on the Wilson Doctrine, such as this paper, it can only seem self-inconsistent to pursue "successive reduction" in the dependence on common knowledge assumptions on one hand, and continue to be obsessed with the CPA on the other.

6 The Impossibility of Bayesian Foundation

In this section, we shall investigate yet another possible foundation of domiannt strategy mechanisms, namely the Bayesian foundation. The Bayesian foundation can be loosely explained with the following story. Imagine the auctioneer as a Bayesian decision maker. When she needs to choose a mechanism under uncertainty of bidders' beliefs, she forms a subjective belief about bidders' beliefs, and compares different mechanisms by calculating the expected performance with respect to that subjective belief. When we as outside observers observe that this auctioneer chooses a particular mechanism, we can ask whether or not such a choice is consistent with Bayesian rationality; i.e., whether or not such a choice is optimal with respect to *some* subjective belief. If the answer is yes, then we say that such a choice is rationalizable. We can say that dominant strategy mechanisms are rationalizable if they are optimal with respect to *some* subjective beliefs.

The difference between the Bayesian and the maxmin foundations is hence whether the auctioneer is a Bayesian or a maxmin decision maker. Given the predominant role of Bayesian rationality in the literature of mechanism design, it seems even more natural to pursue the Bayesian foundation.

To investigate the possibility of the Bayesian foundation, we only need minimal changes in our setting. Recall that we have already been modelling assumptions about bidders' beliefs as distributions over their types. So all we need to do now is to reinterpret an assumption as a subjective belief of the auctioneer. Similarly, if the auctioneer's estimate of the bidders' valuations is described by ν, then her subjective belief about bidders' beliefs must be a distribution μ over bidders' types that is consistent with ν.

It turns out that the Bayesian foundation is impossible. Indeed, it is not difficult to construct examples where dominant strategy mechanisms cannot be rationalized by *any* subjective beliefs. We shall provide one such example below. As it will be clear from the proof below, this example is robust to perturbations.

In this example, there are two bidders and each has two possible valuations. The distribution of valuations ν is represented in Figure 6.

	$v_1 = 5$	$v_1 = 10$
$v_2 = 4$	1/6	0
$v_2 = 2$	1/3	1/2

Figure 6: The distribution ν.

The optimal dsIC auction is depicted in Figure 7, where we follow the convention in Section 3 and use "$\alpha = i$" as the shorthand for "allocating the object to bidder i.".

	$v_1 = 5$	$v_1 = 10$
$v_2 = 4$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$
$v_2 = 2$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$

Figure 7: The optimal dsIC auction Γ.

It is helpful to pay attention to a few noteworthy aspects of this environment and the optimal dsIC auction. Notice that the valuation of bidder 1 is always higher than that of bidder 2. Nevertheless, the auctioneer chooses to sell to bidder 2 when bidder 1 has low valuation. This is optimal because conditional on bidder 2 having low valuation, the probability that bidder 1 has high valuation is greater than 1/2. This means that it is optimal to exclude the low valuation type of bidder 1 to relax the incentive constraint and sell to the high valuation type at his reservation price. Given this, the auctioneer may as well sell to bidder 2 when bidder 1 has a low valuation. If monotonicity were not a constraint, the auctioneer would choose to sell to bidder 1 when bidder 2 had high valuation. Thus, the monotonicity constraint binds here, and in order to satisfy it, the object is sold to bidder 2 in this case.

Proposition 2 *The optimal dsIC auction Γ depicted in Figure 7 cannot be rationalized by any subjective belief μ of the auctioneer that is consistent with the distribution ν depicted in Figure 6.*

In the remainder of this section we will present the proof of Proposition 2. In Appendix C we prove the following stronger result.

Proposition 3 *For the distribution ν depicted in Figure 6, the optimal BIC revenue is uniformly bounded away from the optimal dsIC revenue regardless of the auctioneer's subjective belief; i.e.,*

$$\inf_{\mu \in \mathcal{M}(\nu)} \sup_{\Gamma \text{ is BIC}} R_\mu(\Gamma) > V^D(\nu).$$

To prove Proposition 2, fix any subjective belief $\mu \in \mathcal{M}(\nu)$ that rationalizes the optimal dsIC auction Γ, we shall prove that there exists an BIC auction that generates higher μ-expected revenue than Γ does. This would contradict the assumption that μ rationalizes Γ and complete the proof.

The proof proceeds by a sequence of lemmas. In each we derive conditions that must be satisfied by a rationalizing subjective belief μ. Finally we show that no subjective belief μ can satisfy them all.

For the purpose of this proof, it suffices to work only with bidder 2's first-order beliefs in order to arrive at a contradiction. So, for notational convenience, we shall summarize bidder 2's belief type τ_2 by a single number: his first-order belief that bidder 1 has high valuation. Specifically, for any type $\omega_2 = (v_2, \tau_2)$ of bidder 2, if $v_2 = 4$, we shall use a to denote $g_2(\tau_2)(v_1 = 10)$; and if $v_2 = 2$, we shall use b to denote $g_2(\tau_2)(v_1 = 10)$. For any (measurable) subset $A \subset [0,1]$, we shall use "$a \in A$" to denote the event $\{\omega_2 = (v_2, \tau_2) : v_2 = 4, f_2(\tau_2)(v_1 = 10) \in A\}$; similarly for the notation "$b \in B \subset [0,1]$."

The first lemma says that, conditional on any μ-non-null subset of low-valuation types of bidder 2, the μ-conditional-probability that bidder 1 having high valuation cannot be too low, otherwise the auctioneer can improve upon Γ by selling to some low-valuation types of bidder 1.[5]

Lemma 2 *For any $x \in (0,1]$ such that $\mu(b = x) = 0$, if $\mu(b < x) > 0$, then $\mu(v_1 = 10 | b < x) \geq 3/8$.*

Proof: Suppose there exists $x \in (0,1]$ such that $\mu(b < x) = \mu(b \leq x) > 0$, and yet $\mu(v_1 = 10 | b < x) < 3/8$. Consider the modified auction $\Gamma(x)$ as depicted in Figure 8.

	$v_1 = 5$	$v_1 = 10$
$a \in [0,1]$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$
$b \geq x$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$
$b < x$	$\alpha = 1, t_1 = 5, t_2 = 0$	$\alpha = 1, t_1 = 5, t_2 = 0$

Figure 8: The modified auction $\Gamma(x)$.

[5] In Lemma 2 (and similarly in Lemmas 3-5), the seemingly redundant requirement of $\mu(b = x) = 0$ is a null-boundary property used only in the proof of Proposition 3.

To see that $\Gamma(x)$ continues to be BIC, note that (i) truth-telling continues to be a dominant strategy of bidder 1, (ii) low-valuation types of bidder 2 always have zero rent regardless of what they announce, and (iii) high-valuation types of bidder 2 would not announce the (newly added) message "$b < x$" as that gives them zero rent.

The only difference between $\Gamma(x)$ and Γ is in the (μ-non-null) event of $b < x$, in which case $\Gamma(x)$ generates μ-expected revenue of $5\mu(v_1 = 5|b < x) + 5\mu(v_1 = 10|b < x) = 5$, whereas Γ only generates μ-expected revenue of $2\mu(v_1 = 5|b < x) + 10\mu(v_1 = 10|b < x) < 2(5/8) + 10(3/8) = 5$, contradicting the assumption that μ rationalizes Γ. ■

The second lemma says that for any low-valuation type of bidder 2 that the auctioneer subjectively perceives as possible, his first-order belief b also cannot be too low, otherwise his belief would be too different from the auctioneer's belief, so much so that the auctioneer can improve upon Γ by betting against him.

Lemma 3 *$\mu(b < 3/13) = 0$.*

Proof: Suppose not. Then pick $x < 3/13$ such that $\mu(b < x) > 0$ and $\mu(b = x) = 0$,[6] and consider the modified auction $\Gamma'(x)$ as depicted in Figure 9.

	$v_1 = 5$	$v_1 = 10$
$a \in [0, 1]$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$
$b \geq x$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$
$b < x$	$\alpha = 0, t_1 = 0, t_2 = -2$	$\alpha = 1, t_1 = 10, t_2 = 2(1 - x)/x$

Figure 9: The modified auction $\Gamma'(x)$.

To see that $\Gamma'(x)$ continues to be BIC, note that (i) truth-telling continues to be a dominant strategy of bidder 1, (ii) low-valuation types of bidder 2 would have strict incentive to announce the (newly added) message "$b < x$" if and only if the resulting rent of $2(1 - b) - [2(1 - x)/x]b = 2(1 - b/x)$ is positive, or equivalently if and only if $b < x$, and (iii) high-valuation types of bidder 2 would not announce the (newly added) message "$b < x$" as that gives them rent of $2(1 - a) - [2(1 - x)/x]a = 2(1 - a/x)$, which is lower than the rent of $2(1 - a)$ if they tell the truth.

The only difference between $\Gamma'(x)$ and Γ is in the (μ-non-null) event of $b < x$, in which

[6] It is always possible to pick such an x, as any distribution over $[0, 1]$ can have at most countably many mass points.

case $\Gamma'(x)$ collects from bidder 2 an μ-expected amount of

$$\begin{aligned}
& (-2)\mu(v_1 = 5|b < x) + [2(1-x)/x]\mu(v_1 = 10|b < x) \\
\geq\ & (-2)(5/8) + [2(1-x)/x](3/8) \\
=\ & 3/(4x) - 2 \\
>\ & [3/4(3/13)] - 2 \\
=\ & 5/4
\end{aligned}$$

(where the first inequality follows from Lemma 2), whereas Γ only collects from bidders 2 an μ-expected amount of $2\mu(v_1 = 5|b < x) \leq 2(5/8) = 5/4$, contradicting the assumption that μ rationalizes Γ. ■

The third lemma says that the first-order belief a of high-valuation types of bidder 2 cannot be too low. Otherwise beliefs held by high- and low-valuation types of bidder 2 would be too different, and this would enable the auctioneer to improve upon Γ by introducing Crémer-McLean-kind of bets to separate these types and relax incentive compatibility constraints.

Lemma 4 $\mu(a < 1/11) = 0$.

Proof: If not then let $y < 1/11$ such that $\mu(a = y) = 0$ and $\mu(a < y) > 0$. Notice that $y < 1/11$ implies $y < 3y/(2y+1) < 3/13$, and hence we can also choose x between $3y/(2y+1)$ and $3/13$ such that $\mu(b = x) = 0$. Consider the modified auction $\Gamma(x, y)$ as depicted in Figure 10.

	$v_1 = 5$	$v_1 = 10$
$a < y$	$\alpha = 1, t_1 = 5, t_2 = -2x(1-y)/(x-y)$	$\alpha = 1, t_1 = 5, t_2 = 2(1-x)(1-y)/(x-y)$
$a \geq y$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$
$b < x$	$\alpha = 1, t_1 = 5, t_2 = -2x(1-y)/(x-y)$	$\alpha = 1, t_1 = 5, t_2 = 2(1-x)(1-y)/(x-y)$
$b \geq x$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$

Figure 10: The modified auction $\Gamma(x, y)$.

To see that $\Gamma(x, y)$ continues to be BIC, note that (*i*) truth-telling continues to be a dominant strategy of bidder 1, (*ii*) low-valuation types of bidder 2 would have strict incentive to announce the (newly added) message "$b < x$" if and only if the resulting rent of $[2x(1-y)/(x-y)](1-b) - [2(1-x)(1-y)/(x-y)]b = 2(1-y)(x-b)/(x-y)$ is positive, or equivalently if and only if $b < x$, and (*iii*) high-valuation types of bidder 2 would have strict incentive to announce the (newly added) message "$a < y$" if and only if the resulting rent of $[2x(1-y)/(x-y)](1-a) - [2(1-x)(1-y)/(x-y)]a = 2(1-y)(x-a)/(x-y)$ is strictly higher than the truth-telling rent of $2(1-a)$, or equivalently if and only if $a < y$.

Since the event of $b < x$ is a μ-null event by Lemma 3, the only real difference between $\Gamma(x, y)$ and Γ is in the (μ-non-null) event of $a < y$, in which case $\Gamma(x, y)$ generates μ-expected revenue of

$$\begin{aligned} & 5 - 2x(1-y)/(x-y) \\ = \; & 5 - 2(x - y + y)(1-y)/(x-y) \\ = \; & 5 - 2(1-y) - 2y(1-y)/(x-y) \\ > \; & 5 - 2(1-y) - 2y(1-y)(2y+1)/[3y - y(2y+1)] \\ = \; & 5 - 2(1-y) - 2y(1-y)(2y+1)/[2y(1-y)] \\ = \; & 2, \end{aligned}$$

whereas Γ only generates μ-expected revenue of 2, contradicting the assumption that μ rationalizes Γ. ■

Finally, the fourth lemma says that the first-order belief a of high-valuation types of bidder 2 cannot be too high. Otherwise the beliefs of such types would be too different from the auctioneer's subjective belief, and this would enable the auctioneer to profit by offering an incentive compatible and individually rational bet. Obviously lemmas 4 and 5 deliver the contradiction and thus prove Proposition 2.

Lemma 5 $\mu(a < 1/11) > 0$.

Proof: Suppose $\mu(a < 1/11) = 0$. Consider the modified auction Γ' as depicted in Figure 11.

	$v_1 = 5$	$v_1 = 10$
$a \geq 1/12$	$\alpha = 2, t_1 = 0, t_2 = 123/61$	$\alpha = 2, t_1 = 0, t_2 = 233/61$
$a < 1/12$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$
$b \in [0, 1]$	$\alpha = 2, t_1 = 0, t_2 = 2$	$\alpha = 1, t_1 = 10, t_2 = 0$

Figure 11: The modified auction Γ'.

To see that Γ' continues to be BIC, note that (i) truth-telling continues to be a dominant strategy of bidder 1, (ii) low-valuation types of bidder 2 would not announce the (newly added) message "$a \geq 1/12$" as that gives them strictly negative rent regardless of what bidder 1 announces, and (iii) high-valuation types of bidder 2 would have weak incentive to announce the (newly added) message "$a \geq 1/12$" if and only if the resulting rent of $(4 - 123/61)(1-a) + (4 - 233/61)a$ is weakly higher than their original rent of $2(1-a)$, or equivalently if and only if $a \geq 1/12$.

Since the event $a < 1/12 < 1/11$ is a μ-null event by assumption, the only real difference between Γ' and Γ is in the (μ-non-null) event of $a \geq 1/12$, in which case Γ' generates μ-

expected revenue of $123/61 > 2$, whereas Γ only generates μ-expected revenue of 2. This proves that μ does not rationalize Γ. ∎

7 Conclusion

To summarize: the origin of this study dated back to a paper by Wilson (1987), where he called for "successive reduction" in dependence on common knowledge assumptions. Recent literature responses to his call with the proposal of using stronger solution concepts such as dominant strategy incentive compatibility. Suppose a student of mechanism design is considering following this proposed procedure, but is worried about whether it has any foundation or not. Our suggestion to this student would be: "Follow the proposed procedure first, and check whether the monotonicity constraint is binding or not." If the constraint is not binding, then: "Congratulations. The procedure you just followed does have a foundation."

As we emphasized in the Introduction, we do not think there should be one and only one foundation of dominant strategy mechanisms. Quite the contrary, we foresee the journey of implementing the Wilson Doctrine as successive accumulation of various theories that support the proposed procedure of using dominant strategy mechanisms.

In line with this vision, we shall close this paper by briefly discuss other foundations of dominant strategy mechanisms. Actually, to our knowledge, Bergemann and Morris (2002) is the only other foundation for dominant strategy mechanisms so far. They provide a foundations for ex post incentive compatibility in interdependent-value settings, which nest as a special case dominant strategy incentive compatibility in private-value settings. Our paper and theirs focus on two different subsets of mechanism design problems, and hence these two papers do not overlap with each other. While we focus on optimal auction design problems, they focus on implementation of social choice rules that are measurable with respect to payoff-relevant types.

The reason why we were not able to extend their insight into optimal auction design problems can be loosely explained as follows. In private-value auctions, payoff-relevant types would correspond to bidders' valuations, and social choice rules that are measurable with respect to payoff-relevant types would correspond to allocation rules that are constant over bidders' beliefs. Since many conceivable auction mechanisms would have the property that the way the object is allocated sensitively depends on how bidders' choose among a menu of lotteries, these auction mechanisms would all have allocation rules that are not constant over bidders' beliefs. Examples include the mechanisms depicted in Figures 8, 9, 10, 11, 12, and 13. Had we ruled out these auction mechanisms, we would have in effect "assumed away" the availability of complicated auction mechanisms by brute force.

Appendix A: An Example for Section 5

In Section 5, we claim that there exists a distribution ν that satisfies Condition M such that there is no CPA-assumption μ for which equation 3 holds. We shall provide an example of such a distribution here.

As the proof below would make it clear, this example of ν is a robust to perturbations.

Consider the same example as in Section 3, where there are bidders, and each bidder has two possible valuations. The joint distribution of valuations is depicted in Figure 3, and the corresponding optimal dsIC auction is depicted in Figure 4.

Suppose there exists an CPA-assumption $\mu \in \mathcal{M}(\nu)$ for which equation 3 holds. We shall prove that there exists an BIC auction that generates higher μ-expected revenue than Γ does. This would contradict the supposition that equation 3 holds.

It suffices to work only with bidder 2's first order beliefs in order to complete this proof. So, following the convention in Section 6, we shall continue to use a (b) to denote the first-order belief of a high-valuation (low-valuation) type of bidder 2 that bidder 1 has high valuation. Let $\underline{b} = \sup\{x \in [0,1] : \mu(b < x) = 0\}$.

First, observe that $\underline{b} \geq 4/9$. Suppose, on the contrary, $\underline{b} < 4/9$. Then pick any number z between $\underline{b}$ and $4/9$, and consider the modified auction $\Gamma(z)$ as depicted in Figure 12.

	$v_1 = 4$	$v_1 = 9$
$a \in [0,1]$	$\alpha = 2, t_1 = 0, t_2 = 11$	$\alpha = 2, t_1 = 0, t_2 = 11$
$b \geq z$	$\alpha = 0, t_1 = 0, t_2 = 0$	$\alpha = 1, t_1 = 9, t_2 = 0$
$b < z$	$\alpha = 1, t_1 = 4, t_2 = 0$	$\alpha = 1, t_1 = 4, t_2 = 0$

Figure 12: The modified auction $\Gamma(z)$.

It is obvious that $\Gamma(z)$ continues to be BIC. The only difference between $\Gamma(z)$ and Γ is in the (μ-non-null) event of $b < z$, in which case $\Gamma(z)$ generates μ-expected revenue of 4, whereas Γ only generates μ-expected revenue of $9\mu(v_1 = 9|b < z) < 9z < 9(4/9) = 4$, where the first inequality comes from the fact that μ is an CPA-assumption. Since this would have contradicted the supposition that equation 3 holds, we must have $\underline{b} \geq 4/9$.

Then, consider the modified auction Γ'' as depicted in Figure 13.

To see that Γ'' continues to be BIC, it suffices to observe that, for low-valuation types of bidder 2 with $b \geq 4/9$, truth-telling gives them a non-negative rent of $(5 - 11)(1 - b) + (15/2)b \geq (-6)(5/9) + (15/2)(4/9) = 0$.

Since $b < 4/9$ is a μ-null event, Γ'' generates μ-expected revenue of $9(4/10) + 11(6/10) - (15/2)(4/10) = 72/10$, whereas Γ only generates μ-expected revenue of $9(3/10) + 11(4/10) = 71/10$. This prves that equation 3 does not hold, a contradiction. ■

	$v_1 = 4$	$v_1 = 9$
$a \in [0,1]$	$\alpha = 2, t_1 = 0, t_2 = 11$	$\alpha = 1, t_1 = 9, t_2 = -15/2$
$b \geq 4/9$	$\alpha = 2, t_1 = 0, t_2 = 11$	$\alpha = 1, t_1 = 9, t_2 = -15/2$
$b < 4/9$	$\alpha = 0, t_1 = 0, t_2 = 0$	$\alpha = 0, t_1 = 0, t_2 = 0$

Figure 13: The modified auction Γ''.

Appendix B: Proof of Proposition 1

Proof of Proposition 1 Let $\gamma_i(v)$ denote the derivative of (5) with respect to $p_i(v)$ multiplied by $1/\nu(v)$. We have

$$\gamma_i(v) = v_i - \Delta h_i(v).$$

The value $\gamma_i(v)$ is the dsIC-analogue of virtual utility for bidder i at the valuation profile v. If the monotonicity constraint were not present, the objective function in (5) would be maximized by assigning the object to any bidder with the highest non-negative virtual utility and withholding the object when all virtual utilities were negative. Let p be such an allocation rule. Condition M is satisfied if p is monotonic. Suppose $h_i(v_i, v_{-i}) - h_j(v_i, v_{-i})$ is non-increasing in v_i. We argue that in this case p is monotonic. For if $p_i(v^m, v_{-i}) > 0$, then i has (among the) highest virtual utility at (v^m, v_{-i}). The non-decreasing inverse hazard-rate condition implies that i must have the strictly highest virtual utility at (v^{m+1}, v_{-i}) and hence that $p_i(v^{m+1}, v_{-i}) = 1$. It follows that p is monotonic. ■

Appendix C: Proof of Proposition 3

Lemma 6 *Suppose K is a compact topological space and that $\mathcal{F}$ is a family of real-valued functions on K such that, for each $x \in K$, there is some $f_x \in \mathcal{F}$ which is continuous at x and satisfies $f_x(x) > 0$. Then we have $\inf_{x \in K} \sup_{f \in \mathcal{F}} f(z) > 0$.*

Proof: For each $x \in K$, there exists an open neighborhood U_x such that, for each $y \in U_x$, we have $f_x(y) > f_x(x)/2$. The collection $\{U_x : x \in K\}$ forms an open covering of the compact space K, and hence there exists a finite subcovering. Let $\{U_{x_1}, \ldots, U_{x_n}\}$ be a finite subcovering and let $\varepsilon = \min\{f_{x_1}(x_1), \ldots, f_{x_k}(x_n)\} > 0$. For each $x \in K$, we have $x \in U_{x_l}$ for some $l = 1, \ldots, n$ so that $\sup_{f \in \mathcal{F}} f(x) \geq f_{x_l}(x) > f_{x_l}(x_l)/2 \geq \varepsilon/2 > 0$. ■

Lemma 7 *Suppose $\mathcal{O}_1, \ldots, \mathcal{O}_n$ are disjoint open subsets of Ω^* such that $\mu(\cup \mathcal{O}_l) = 1$, and $t : \Omega^* \to \mathbf{R}$ is a bounded real function that is constant on each $\mathcal{O}_l$. Then the mapping*

$$\mu' \to \int_{\Omega^*} t \, \mu'(d\omega)$$

is continuous at the point μ.

Proof: Fix any $\varepsilon > 0$. Let $\bar{t} > 0$ be an upper bound for $|t|$. The function $\mu' \to \mu'(\mathcal{O}_i)$ is lower semi-continuous (see Aliprantis and Border (1999)), hence we can set

$$\delta = \frac{\varepsilon}{\bar{t}n^2}$$

and find a neighborhood U of μ such that, for all $\mu' \in U$, $\mu'(\mathcal{O}_l) > \mu(\mathcal{O}_l) - \delta$ for $l = 1, \ldots, n$. Since $\mu(\cup\mathcal{O}_l) = 1$, it follows that $\mu'(\mathcal{O}_l) < \mu(\mathcal{O}_l) + (n-1)\delta$ and $\mu'(\Omega^* \setminus \cup\mathcal{O}_l) < \mu(\Omega^* \setminus \cup\mathcal{O}_l) + n\delta = n\delta$.

We can write

$$\int_{\Omega^*} t\, d\mu' = \sum_{l=1}^{n} \mu'(\mathcal{O}_l)t(\mathcal{O}_l) + \int_{\Omega^* \setminus \cup\mathcal{O}_l} t(\omega)\, d\mu',$$

so that

$$\begin{aligned}
& \sum_{l=1}^{n} \mu'(\mathcal{O}_l)t(\mathcal{O}_l) - \mu(\Omega^* \setminus \cup\mathcal{O}_l)\bar{t} \leq \int_{\Omega^*} t\, \mu'(d\omega) \leq \sum_{l=1}^{n} \mu'(\mathcal{O}_l)t(\mathcal{O}_l) + \mu'(\Omega^* \setminus \cup\mathcal{O}_l)\bar{t} \\
\Longrightarrow \quad & \sum_{l=1}^{n}[\mu(\mathcal{O}_l) - \delta]t(\mathcal{O}_l) - n\delta\bar{t} < \int_{\Omega^*} t\, \mu'(d\omega) < \sum_{l=1}^{n}[\mu(\mathcal{O}_l) + (n-1)\delta]t(\mathcal{O}_l) + n\delta\bar{t} \\
\Longrightarrow \quad & -\delta \sum_{l=1}^{n} t(\mathcal{O}_l) - n\delta\bar{t} < \int_{\Omega^*} t\, \mu'(d\omega) - \int_{\Omega^*} t\, \mu(d\omega) < (n-1)\delta \sum_{l=1}^{n} t(\mathcal{O}_l) + n\delta\bar{t} \\
\Longrightarrow \quad & -2n\delta\bar{t} < \int_{\Omega^*} t\, \mu'(d\omega) - \int_{\Omega^*} t\, \mu(d\omega) < n^2\delta\bar{t}.
\end{aligned}$$

This proves that $\left|\int_{\Omega^*} t\, \mu'(d\omega) - \int_{\Omega^*} t\, \mu(d\omega)\right| < \max\{2n\delta\bar{t}, n^2\delta\bar{t}\} = \varepsilon$. ■

Proof of Proposition 3 Notice that, for each of the mechanisms used in the proof of Proposition 2, the total transfer $(t_1+t_2)(\omega)$ satisfies the conditions of Lemma 7. For example, consider the mechanism $\Gamma(x)$ in Lemma 2. For any (v_1, v_2), the set of universal type profiles in which the valuation pair is (v_1, v_2) is open in the product topology with μ-null boundary. Moreover, since $\mu(b = x) = 0$, the event $b < x$ is also open in the product topology with μ-null boundary. Therefore, we can take $\mathcal{O}_1, \ldots, \mathcal{O}_6$ to be the interiors of the sets represented by the cells of the table in Figure 8. These open sets are disjoint, have μ-null boundaries, and have total μ-measure equal to 1 as required.

Thus, for any auctioneer's belief μ that is consistent with the distribution ν, there exists an BIC auction $\Gamma(\mu)$ such that $R_\mu\Gamma(\mu) - V^D(\nu) > 0$, and the mapping $\mu' \to R_{\mu'}\Gamma(\mu) - V^D(\nu)$ is continuous at the point $\mu' = \mu$. We can hence apply Lemma 6, taking $K = \mathcal{M}(\nu)$ and $\mathcal{F} = \{R_{(\cdot)}\Gamma - V^D(\nu) : \Gamma \text{ is BIC}\}$. ■

References

ALIPRANTIS, C. D., AND K. C. BORDER (1999): *Infinite Dimension Analysis: A Hitchhiker's Guide*. Springer-Verlag, New York.

BERGEMANN, D., AND S. MORRIS (2002): "Robust Mechanism Design," mimeo, Yale University.

BRANDENBURGER, A., AND E. DEKEL (1993): "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189–198.

CRÉMER, J., AND R. P. MCLEAN (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257.

DASGUPTA, P., AND E. MASKIN (2000): "Efficient Auctions," *Quarterly Journal of Economics*, 115(2), 341–388.

GUL, F. (1998): "A Comment on Aumann's Bayesian View," *Econometrica*, 66, 923–928.

LIPMAN, B. (2003): "Finite Order Implications of the Common Prior Assumption," *Econometrica*.

MERTENS, J. F., AND S. ZAMIR (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1–29.

MORRIS, S. (1995): "The Common Prior Assumption in Economic Theory," *Economics and Philosophy*, 11, 227–253.

MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

PERRY, M., AND P. J. RENY (2002): "An Efficient Auction," *Econometrica*, 70(3), 1199–1212.

WILSON, R. (1987): "Game-Theoretic Approaches to Trading Processes," in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley. Cambridge University Press.