האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

ON THE GENERIC (IM)POSSIBILITY OF FULL SURPLUS EXTRACTION IN MECHANISM DESIGN

by

AVIAD HEIFETZ and ZVIKA NEEMAN

Discussion Paper # 350   March 2004

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135   FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

On the Generic (Im)possibility of Full Surplus Extraction in Mechanism Design

Aviad Heifetz* Zvika Neeman†

Preliminary Version: February 27, 2004

Abstract

A number of studies, most notably Crémer and McLean (1985, 1988), have shown that in Harsanyi type spaces of a fixed finite size, it is generically possible to design mechanisms that extract all the surplus from players, and as a consequence, implement any outcome as if the players' private information were commonly known. In contrast, we show that within the set of common priors on the universal type space, the subset of priors that permit the extraction of the players' full surplus is *shy.* Shyness is a notion of smallness for convex subsets of infinite-dimensional topological vector spaces (in our case, the set of common priors), which generalizes the usual notion of zero Lebesgue measure in finite-dimensional spaces.

Journal of Economic Literature Classification numbers: D44, D82, H41.

Keywords: Surplus extraction, information rents, mechanism design, private information, universal type space, genericity, prevalence.

*Department of Economics and Management, The Open University of Israel, aviadhe@openu.ac.il

†Department of Economics, Boston University, and Center for Rationality and Department of Economics, Hebrew University of Jerusalem, http://people.bu.edu/zvika, zvika@bu.edu

1 Introduction

Does the holding of relevant private information necessarily confers a positive economic rent? Surprisingly, the answer given by the literature to this question is negative. A number of studies, including, most notably, Crémer and McLean (1985, 1988), have shown that under standard assumptions such as the existence of a common prior, a fixed finite number of types, risk neutrality, and no limited liability, it is generically possible to implement any outcome as if the players' private information were commonly known. In particular, a seller, for example, should generically be able to extract the full surplus of any number of bidders in an auction. As these "full-surplus-extraction" results imply that the players' private information is (generically) irrelevant, they have been said to "cast doubt on the value of the current mechanism design paradigm as a model of institutional design" (McAfee and Reny, 1992, p. 400).

Since full-surplus-extraction results make heavy use of the fixed finite type space assumption, it is natural to ask whether or not the possibility of full-surplus-extraction extends to the most general consistent private information type space imaginable, or to the collection of consistent (in the sense of Harsanyi, 1967-68) subspaces of the universal type space (Mertens and Zamir, 1985), in which it is common knowledge that each player knows her own signal, and in which every other consistent private information type space can be embedded. Supposing that any common prior on this universal type space could just as well serve as a plausible model of a situation involving asymmetric information, is it "typically" the case that full-surplus-extraction is possible? This is the question addressed in this paper.[1] The assumption of risk neutrality and no limited liability which is the other important assumption necessary for full surplus extraction is maintained throughout this paper.

If beliefs are endowed with the minimal topology that allows for the formulation of each player's beliefs about the state of nature and the beliefs of other players (the topology of weak convergence), then the set of common priors with finite support is dense in the space of all common priors (Mertens and Zamir, 1985). Combining this observation with the results of Crémer and McLean (1985, 1988) for finite type spaces implies that the set of priors that permit full-surplus-extraction is dense in the space of all possible common priors.

Recently, Neeman (2001) showed that full-surplus-extraction is possible only if the type space has the property that every possible belief of every player about other players' types is associated (with probability one) with a unique valuation or private signal of the player (Neeman called this property "beliefs determine preferences"). Neeman (2001) also showed that arbitrarily close to any consistent finite type space, there is another consistent finite type

[1] The reason we confine our attention to consistent subspaces of the universal type space is twofold. First, confining attention to Harsanyi-consistent subspaces of the universal type space and their associated common priors is standard practice in information economics. Indeed, this practice, which has been called the Harsanyi doctrine, is considered by some to be a hallmark of Bayesian rationality (see the discussion in Aumann 1998, Gul 1998 and Morris 1995). Second, the universal space has a product structure – it includes all the possible combinations of players' beliefs and private signals. Hence, as should become clear below when the terms are defined, beliefs do not determine preferences in the universal type space, and consequently full-surplus-extraction is generically impossible there. Thus the question is settled in this case.

space in which beliefs do not determine preferences, and consequently full surplus cannot be extracted. It thus follows that both the set of priors that allow for full-surplus-extraction (henceforth, FSE priors), *and* its complement – the set of priors that do not allow for full-surplus-extraction (henceforth, NFSE priors), are *dense* in the space of all finite-support priors, and hence also in the space of all priors. In particular, with the topology of weak convergence, there is no topological sense in which one of these sets can be said to be larger (i.e., open and dense) than the other.

However, just as both the rationals and the irrationals are dense in the set of real numbers although the set of irrationals is larger in other senses (cardinality, Lebesgue measure), it is also conceivable that one of the subsets of common priors above is larger than the other in some meaningful sense.

One may think of two general approaches that may permit such a sharper result. First, it may be argued that the topology of weak convergence on priors is not the natural topology to apply in a strategic setting.[2] Intuitively, two priors can be said to be "close" if and only if they induce "similar" equilibrium behavior. Indeed, preliminary results (Kajii and Morris, 1998) suggest that such a notion of strategic proximity may induce a stronger topology on the set of priors. However, a proper definition of strategic proximity in general type spaces and its characterization in terms of beliefs is not yet available. If and when such a characterization is obtained, then it may turn out that the resulting stronger topology renders one of the subsets of common priors above both open and dense, while its complement would not be dense.

Second, it is also possible to consider non-topological notions of genericity, such as the measure-theoretic notion of full Lebesgue measure. Unfortunately, the notion of full Lebesgue measure cannot be applied directly because it can meaningfully capture the idea of a set being "large" only in finite-dimensional spaces. In contrast, the space of common priors on the universal type space (or even the smaller space of common priors with a finite support) is not only infinite but infinite-dimensional. It is therefore necessary to consider a measure-theoretic notion of genericity that can be applied in infinite-dimensional topological vector spaces. Such an appropriate notion, called *prevalence,* was originally conceived by Christensen (1974) and Hunt et al. (1992) and further developed by Anderson and Zame (2001) as an extension of the idea of full Lebesgue measure to infinite-dimensional spaces. The complement of a prevalent set is called *shy.* A collection of shy sets in a finite-dimensional space is identical to a collection of sets with Lebesgue measure zero. In an infinite-dimensional space, shy sets retain the properties of zero-probability events: no open set is shy, and the collection of shy sets is closed under subsets, under translations, and under countable unions.

We show that the subset of FSE priors is shy within the set of common priors on the universal type space. It therefore follows that the complement of the set of FSE priors, or the subset of NFSE priors, is generic. The same result also obtains if attention is restricted to the subset of priors with finite-support.

The proof is based on the following lemma, which stems from Neeman's (2001) observation that full-surplus-extraction requires that players' beliefs determine their private signals

[2] See Morris (2002) for such an argument.

or preferences. Only weighted averages of FSE priors yield FSE priors, while a weighted average of a NFSE and any other prior yields a NFSE prior. This asymmetry "in favor" of the NFSE priors delivers the result. What makes this result mathematically non-trivial is the fact that the set of FSE priors is dense in the consistent universal type space.

The rest of the paper proceeds as follows. For simplicity, instead of considering a general mechanism design problem with interdependent players' types, we begin in the next section with the consideration of the classic problem of a seller of an object who designs an auction for n risk neutral bidders with private valuations, with the goal of maximizing his expected revenue. In Section 3, we explain how our results can be applied to any mechanism design problem with interdependent types. Section 4 surveys the related literature, and contains a discussion of the relationship of our results to those of Crémer and McLean.

2 Surplus Extraction in Single Object Auctions with Private Values

We consider the problem of a seller who wishes to design an auction that would maximize the expected revenue he obtains by selling some object to one out of a set of n risk neutral bidders with private valuations for the object. Each bidder or player may refuse to participate in the seller's auction, but if she agrees to participate, then she is bound by the outcome of the auction.

Let $N = \{1, ..., n\}$ denote the set of bidders or players. Each bidder i has a certain nonnegative valuation or willingness to pay for the object which we denote by $v_i \in V_i$. The set of bidder i's valuations V_i is assumed to be a complete, separable, metric space (in particular, V_i may be finite). The payoff to a bidder with valuation v_i who wins the object with probability q and who pays an expected amount x is given by $q \cdot v_i - m$. We refer to v_i as bidder i's *preference* or *preference type*. Let $V \equiv V_1 \times \cdots \times V_n$. The set V is the *basic space of uncertainty* for this problem.

The behavior of the bidders in the auction may obviously depend on their willingness to pay for the object. It therefore follows that the bidders' behavior may also depend on their beliefs about other bidders' willingness to pay, because such beliefs convey possibly important information about the way other bidders will behave in the auction. For the same reason, beliefs about beliefs are also important, and so are beliefs about beliefs about beliefs, and so on, ad infinitum. A complete analysis of the seller's problem therefore requires a model that allows for the specification of the bidders' entire infinite hierarchy of beliefs about beliefs about beliefs ... about whatever is relevant in the auction. Such infinite hierarchies of beliefs may be conveniently encoded in what is known as a *type space*.

2.1 Type Spaces

Bidder i's private information is captured by its type $\theta_i \in \Theta_i$. The sets of bidders' types Θ_i, $i \in N$, are assumed to be complete, separable, metric spaces. For every space X, let $\Delta(X)$ denote the space of probability measures over X. Each type $\theta_i \in \Theta_i$ is associated with a *preference type* $\widehat{v}_i(\theta_i) \in V_i$ which describes θ_i's willingness to pay for the object, and with a *belief-type* $\widehat{b}_i(\theta_i) \in \Delta(\Theta_{-i})$ which is a belief, or a probability measure, on the space of other bidders' types $\Theta_{-i} \equiv \prod_{j \neq i} \Theta_j$. The space of probability measures $\Delta(\Theta_{-i})$ is endowed with the topology of weak convergence.

Each type of each bidder is assumed to know its own willingness to pay for the object and its beliefs. Because we focus our attention in this section on a private values model, each type θ_i's preference type $\widehat{v}_i(\theta_i)$ is defined independently of θ_i's belief type $\widehat{b}_i(\theta_i)$. This assumption is relaxed in the next section.[3]

A product space $\Theta \equiv \prod_{i \in N} \Theta_i$ of the players' type spaces is called a private values type space. Each profile of types $\theta \in \Theta$ is called a state of the world.

2.2 The Private Values Universal Type Space

Given the basic space of uncertainty $V \equiv V_1 \times \cdots \times V_n$ and the set of bidders N, there exists[4] a *private values universal type space*

$$T^{PV} = \prod_{i \in N} T_i^{PV}$$

into which every other private values type space can be mapped in a beliefs-preserving way. That is, for every type space $\Theta \equiv \prod_{i \in N} \Theta_i$ there exists a unique set of measurable mappings[5]

$$\mathbb{E}_i : \Theta_i \to T_i^{PV}, \quad i \in N$$

satisfying

$$\widehat{v}_i(\mathbb{E}_i(\theta_i)) = \widehat{v}_i(\theta_i)$$

and

$$\widehat{b}_i(\mathbb{E}_i(\theta_i))(A) = \widehat{b}_i(\theta_i)\left(\mathbb{E}_{-i}^{-1}(A)\right)$$

for every measurable set $A \subseteq T_{-i}^{PV}$, where

$$\mathbb{E}_{-i} : \Theta_{-i} \to T_{-i}^{PV}$$

is defined by

$$\mathbb{E}_{-i}\left((\theta_j)_{j \neq i}\right) = (\mathbb{E}_j(\theta_j))_{j \neq i}.$$

The universal type set T_i^{PV} of bidder $i \in N$ is isomorphic to the product space

$$V_i \times \Delta\left(T_{-i}^{PV}\right)$$

[3] The assumption that each type knows its own belief is captured by defining $\widehat{b}_i(\theta_i)$ as a probability measure over Θ_{-i} rather than over $\Theta_i \times \Theta_{-i}$. The implied presumption about the bidders' introspective ability is standard, and is maintained throughout the paper.

[4] The proof of existence follows from a slight adaptation of the arguments in Mertens and Zamir (1985), Brandenburger and Dekel (1993), and Heifetz (1993).

[5] which are in fact also continuous

by the mapping

$$\tau_i \to \left(\widehat{v}_i\left(\tau_i\right), \widehat{b}_i\left(\tau_i\right)\right).$$

Thus in what follows we use the terms T_i^{PV} and $V_i \times \Delta\left(T_{-i}^{PV}\right)$ interchangeably.

Finally, for the rest of this section we drop the superscript "PV" from the notation whenever there is no risk of confusion.

2.3 Priors

A probability measure p_i on a private values type space $\Theta = \prod_{i \in N} \Theta_i$ is called a prior for bidder i if bidder i's belief-types $\widehat{b}_i\left(\theta_i\right)$ are the posteriors of p_i. That is, p_i is a prior for bidder i if for every real-valued continuous function $\varphi : \Theta \to \mathbb{R}$

$$\int_{\Theta_i} \left(\int_{\Theta_{-i}} \varphi\left(\theta_i, \tilde{\theta}_{-i}\right) d\widehat{b}_i\left(\theta_i\right)\left(\tilde{\theta}_{-i}\right)\right) dp_{i|\Theta_i}\left(\theta_i\right) = \int_{\Theta} \varphi\left(\theta\right) dp_i\left(\theta\right) \tag{1}$$

where $p_{i|\Theta_i}$ is the marginal of p_i on Θ_i.

A probability measure p on Θ is called a common prior, or prior for short, if it is a prior for every bidder $i \in N$. The support of a prior p is called a Harsanyi-consistent subspace.

It is immediate from the definition that the set of priors for bidder i is convex: If $p_i, p_i' \in \Delta\left(\Theta\right)$ are two priors for bidder i, then so is $\alpha p_i + \left(1-\alpha\right) p_i'$ for every $\alpha \in [0,1]$. It follows that the set of common priors is also convex.

Because every private values type space can be embedded in the private values universal type space, no loss of generality is implied by restricting attention to priors on the private values universal type space T. We thus take the set of priors on the private values universal space, denoted $\mathcal{P}^{PV}$, to be the set of relevant "environments" for our study.

Definition. A prior $p \in \Delta\left(T\right)$ satisfies the *Beliefs-Determine-Preferences* (BDP) property[6] for bidder $i \in N$ if there exists a measurable subset $T_i^p \subseteq T_i$ such that the marginal $p_{|T_i}$ of p on T_i assigns probability 1 to T_i^p,

$$p_{|T_i}\left(T_i^p\right) = 1,$$

and no pair of distinct types $\tau_i \neq \tau_i'$ in T_i^p hold the same beliefs –

$$\hat{b}\left(\tau_i\right) \neq \hat{b}\left(\tau_i'\right)$$

for every two different types $\tau_i, \tau_i' \in T_i^p$.

A prior p that satisfies the beliefs-determine-preferences property for bidder i is called a BDP prior for bidder i. A prior p that is a BDP prior for every bidder $i \in N$ is called a BDP prior. Because any pair of distinct types $\tau_i \neq \tau_i'$ in the private values universal space differ

[6] The notion of beliefs-determine-preferences generalizes the one in Neeman (2001) and is closely related to Bergemann and Morris (2003) one-to-one property and to d'Aspremont et al.'s (2002) notion of no free beliefs.

either by their belief-type or by their preference-type,[7] there is no pair of distinct types in T_i^p who hold identical beliefs but different preferences. In other words, a prior p satisfies the BDP property for bidder i if there exists a p-probability 1 set $V^i \times B_i^p$ where $B_i^p \subseteq \Delta(T_{-i})$, and a function that maps bidder i's beliefs to its willingness to pay

$$\Phi_i^p : B_i^p \to V_i$$

such that T_i^p is isomorphic to the graph $\left\{ \left(\widehat{b}_i(\tau_i), \widehat{v}_i(\tau_i) \right) : \tau_i \in T_i^p \right\}$ of Φ_i^p.

We show that BDP is necessary for full-surplus-extraction. Specifically, we show that if a prior p permits the extraction of bidder i's full surplus p-almost surely, then p is a BDP prior for player i.

By the revelation principle, no loss of generality is implied by assuming that the seller employs an incentive compatible and individually rational "direct revelation" auction game or mechanism $\langle q_i : T \to [0,1], m_i : T \to [0,1] \rangle_{i \in N}$ in which each bidder i is asked to report its type $\tau_i \in T_i$, and then to participate in a lottery in which it pays an amount $m_i(t)$, and wins the object with probability $q_i(t)$.

A mechanism $\langle q_i, m_i \rangle_{i \in N}$ is incentive-compatible (IC) if every type $\tau_i \in T_i$ of every bidder $i \in N$ maximizes its expected payoff by truthfully reporting its type, or

$$\int_{T_{-i}} \left(q_i(\tau_i, \tilde{\tau}_{-i}) \hat{v}_i(\tau_i) - m_i(\tau_i, \tilde{\tau}_{-i}) \right) d\hat{b}_i(\tau_i)(\tilde{\tau}_{-i})$$
$$\geq \int_{T_{-i}} \left(q_i(\tau_i', \tilde{\tau}_{-i}) \hat{v}_i(\tau_i) - m_i(\tau_i', \tilde{\tau}_{-i}) \right) d\hat{b}_i(\tau_i)(\tilde{\tau}_{-i}) \quad \text{(IC)}$$

for every $\tau_i' \in T_i$.

A mechanism $\langle q_i, m_i \rangle_{i \in N}$ is individually-rational if every type $\tau_i \in T_i$ of every bidder $i \in N$ prefers to participate in the mechanism than to opt out, or

$$\int_{T_{-i}} \left(q_i(\tau_i, \tilde{\tau}_{-i}) \hat{v}_i(\tau_i) - m_i(\tau_i, \tilde{\tau}_{-i}) \right) d\hat{b}_i(\tau_i)(\tilde{\tau}_{-i}) \geq 0. \quad \text{(IR)}$$

Definition. A prior p permits the *full-surplus-extraction* from a set $K \subseteq N$ of bidders if there exists an incentive compatible and individually rational mechanism $\langle q_i, m_i \rangle_{i \in N}$ that generates an expected payment to the seller that is equal to the full surplus generated by the bidders in K or

$$\sum_{i \in K} \int_T m_i(\tau)\, dp(\tau) = \int_T \max_{i \in K} \{\hat{v}_i(\tau_i)\}\, dp(\tau).$$

A prior that permits the full-surplus-extraction from the set K of bidders is called a full-surplus-extraction (FSE) prior for K.

[7] If there were two distinct types $\tau_i \neq \tau_i'$ with the same preferences and beliefs in the private values universal type space, then the beliefs-preserving mappings of type spaces into the universal space would not have been unique, in contradiction to the definition of the universal type space.

Remark 1. Fix a prior p. In order to extract the full surplus from the bidders in $K \subseteq N$ in the environment described by p, in p-almost surely every state of the world $\tau \in T$, the seller must sell the object to the bidder in K who has the highest willingness to pay for it at τ at an expected price that is equal to this bidder's valuation of the object. The individual rationality constraint implies that the seller cannot sell the object to any bidder in K at τ for a higher price, and selling the object for a lower price implies a failure to extract the full surplus that is generated by the bidders in K. It therefore follows that the individual rationality constraint of any type $\tau_j \in T_j$ of a bidder in K that wins the object with a positive probability under a mechanism $\langle q_i, m_i \rangle_{i \in N}$ that extracts the full surplus of the bidders in K must be binding.

Remark 2. Note that the fact that a prior may permit the extraction of the full surplus from the set N of bidders does not imply that it is possible to extract the full surplus of each single bidder. For example, if there are two bidders and it is commonly known that bidder 1's willingness to pay for the object is strictly lower than that of player 2's, it is possible to extract the full surplus of the set of bidders $\{1, 2\}$ if and only if it is possible to extract the full surplus of bidder 2 alone. Conversely, it can be shown that if it is possible to extract the full surplus of each single bidder, then it is also possible to extract the full surplus of the set N of bidders.

Proposition 1. A prior p that is a FSE prior for bidder i is a BDP prior for bidder i.

Proof. Suppose that p is a FSE prior for bidder i. Let $\langle q_i, m_i \rangle_{i \in N}$ be an incentive compatible and individually rational mechanism that extracts the full surplus of bidder i. By Remark 1, bidder i must win the object with p-probability 1 under the mechanism $\langle q_i, m_i \rangle_{i \in N}$, and bidder i's individual rationality constraint must be binding with p-probability 1 under the mechanism $\langle q_i, m_i \rangle_{i \in N}$.

Suppose that p is not a BDP prior for bidder i. It follows that there exist two disjoint measurable subsets of bidder i's types, $A_i, A'_i \subseteq T_i$, that each have a positive p-probability

$$p_{|T_i}(A_i) > 0, \quad p_{|T_i}(A'_i) > 0,$$

and the same range of beliefs

$$\hat{b}_i(A_i) = \hat{b}_i(A'_i) \subseteq \Delta(T_{-i}),$$

but different valuations. That is, if $\tau_i \in A_i$ and $\tau'_i \in A'_i$ are such that

$$\hat{b}_i(\tau'_i) = \hat{b}_i(\tau_i)$$

then

$$\hat{v}_i(\tau'_i) < \hat{v}_i(\tau_i).$$

In particular, for every type $\tau_i \in A_i$ there exists a type $\tau_i' \in A_i'$ such that $\hat{b}_i(\tau_i') = \hat{b}_i(\tau_i)$ but $\hat{v}_i(\tau_i') < \hat{v}_i(\tau_i)$. It follows that

$$\begin{aligned}
& \int_{T_{-i}} \left(q_i(\tau_i, \tilde{\tau}_{-i})\, \hat{v}_i(\tau_i) - m_i(\tau_i, \tilde{\tau}_{-i})\right) d\hat{b}_i(\tau_i)(\tilde{\tau}_{-i}) \\
\geq & \int_{T_{-i}} \left(q_i(\tau_i', \tilde{\tau}_{-i})\, \hat{v}_i(\tau_i) - m_i(\tau_i', \tilde{\tau}_{-i})\right) d\hat{b}_i(\tau_i)(\tilde{\tau}_{-i}) \\
= & \int_{T_{-i}} \left(q_i(\tau_i', \tilde{\tau}_{-i})\, \hat{v}_i(\tau_i) - m_i(\tau_i', \tilde{\tau}_{-i})\right) d\hat{b}_i(\tau_i')(\tilde{\tau}_{-i}) \\
> & \int_{T_{-i}} \left(q_i(\tau_i', \tilde{\tau}_{-i})\, \hat{v}_i(\tau_i') - m_i(\tau_i', \tilde{\tau}_{-i})\right) d\hat{b}_i(\tau_i')(\tilde{\tau}_{-i}) \\
\geq & \ 0.
\end{aligned}$$

The first inequality follows from the (IC) constraint for type τ_i; the following equality follows from the fact that $\hat{b}_i(\tau_i') = \hat{b}_i(\tau_i)$; the next strict inequality follows from the fact that $\hat{v}_i(\tau_i') < \hat{v}_i(\tau_i)$ and that $q_i(\tau_i', \tilde{\tau}_{-i}) = 1$ for p-almost every type $\tau_i' \in A_i'$; and the last inequality follows from the (IR) constraint for type τ_i'. It therefore follows that bidder i's individual rationality constraint is not binding for p-almost every type $\tau_i \in A_i$. A contradiction. ■

Conversely, it can be shown that if p is a BDP prior for bidder i, then for every $\varepsilon > 0$ it is possible to extract bidder i's full surplus up to ε (McAfee and Reny, 1992); if, in addition, bidder i has only finitely many types, then it is possible to extract bidder i's full surplus (Crémer and McLean, 1988).

2.4 Genericity

In finite dimensional spaces, genericity is often identified with full Lebesgue measure. A set that has Lebesgue measure zero is considered nongeneric, small, or "atypical." A set that has full Lebesgue measure is considered generic, large, or "typical." The situation in infinite-dimensional spaces is more complicated. Unlike the Lebesgue measure in a finite-dimensional Euclidian space $\mathbb{R}^k$, which is spread uniformly across the space, in infinite-dimensional spaces there is no (sigma-additive) measure that "fills up the space." For example, in an infinite-dimensional separable Banach space, any open ball of radius $r > 0$ contains an infinite sequence of disjoint open balls of radius $\frac{r}{4}$, so if a translation-invariant measure were to assign a positive measure to these balls, then the r-ball would be assigned an infinite measure for any $r > 0$.[8] Therefore, probabilities or measures in infinite-dimensional spaces are not satisfactory devices for determining whether events are "typical" or not.

An appealing notion of "smallness" in infinite-dimensional spaces is based on the observation that an event E in a finite-dimensional Euclidian space $\mathbb{R}^k$ has Lebesgue measure zero

[8] Furthermore, confining attention to full-support *quasi-invariant* measures, which preserve null-sets under translations (such as the Gaussian measures on the Euclidean spaces), is unhelpful either. Under fairly general conditions, if there does not exist a non-trivial full support invariant measure on an infinite-dimensional space, then neither does there exist such a quasi-invariant measure (see, e.g., Yamasaki 1985).

if and only if there exists a positive measure μ on $\mathbb{R}^k$ such that E and all its translations $\{x+y : x \in E\}$, $y \in \mathbb{R}^k$, have μ-measure zero. Christensen (1974) and Hunt, Sauer, and York (1992) have relied on this observation to propose a notion of "smallness" that coincides with full Lebesgue measure in finite dimensional spaces and that extends naturally to infinite-dimensional spaces. They defined a Borel subset of a complete metric topological vector space to be *shy* if there exists a positive measure μ on the space such that the set and all its translations have μ-measure zero, and called the complement of a shy set *prevalent*. They showed that shy sets satisfy all the requirements one would expect from "small" or "negligible" events. In particular, a subset of a shy set is shy, every translation of a shy set is shy, a countable union of shy sets is shy, and no open set is shy.

Anderson and Zame (2001) have adapted Christensen (1974) and Hunt et al.'s (1992) definition to the case in which the relevant parameter set is a convex subset C of a topological vector space X. Because we are interested in determining the genericity of the set of FSE priors relative to the space of priors on the universal type space, this is the definition which is appropriate for our purpose.

It turns out that for our analysis it is not necessary to rely on Anderson and Zame's general definition of shyness, but rather on a simpler and stronger notion called "finite shyness." Let λ_H denote the Lebesgue measure on a finite-dimensional subspace $H \subseteq X$.

Definition. (Anderson and Zame, 2001) A universally measurable[9] subset $E \subseteq C$ is *finitely shy* in $C \subseteq X$ if there exists a finite-dimensional subspace $H \subseteq X$ such that $\lambda_H(C+p) > 0$ for some $p \in X$ and $\lambda_H(E+x) = 0$ for every $x \in X$. An arbitrary subset $F \subseteq X$ is finitely shy in C if it is contained in a finitely shy universally measurable set.

Anderson and Zame (2001) show that if a set E is finitely shy in C then it is also shy in C. A subset $Y \subseteq C$ is said to be prevalent in C if its complement $C \backslash Y$ is shy in C.[10]

Example 1. Anderson and Zame (2001). Everywhere differentiable concave functions are finitely shy in the cone of all concave functions.

Example 2. Stinchcombe (2001). Both the subspaces of purely atomic measures and purely non-atomic measures are finitely shy in the space of all measures (in any topology in which they are Borel, or more generally, universally measurable).

[9] A subset $E \subseteq X$ is universally measurable if it is measurable with respect to the completion of every regular Borel probability measure on X.

[10] In their definition, Anderson and Zame required the convex subset $C \subseteq X$ to be completely metrizable, but as they mention in a footnote, the definition makes sense even without this requirement, which is needed only for establishing some enhanced properties of shyness and prevalence (e.g., if E is prevalent in F and F is prevalent in G then E is prevalent in G). This additional requirement is not needed for establishing the basic properties of shy sets, namely that a subset of a shy set is shy, that every translation of a shy set is shy, that a countable union of shy sets is shy, and that no open set is shy.

2.5 FSE Priors are Non-Generic

In this section we show that the set of FSE priors, denoted $\mathcal{F}$, is finitely shy in the set of priors on the private values universal type space, $\mathcal{P}$.

Positive multiples of priors in $\mathcal{P}$ constitute a convex cone of (positive) measures. Taking the differences of pairs of such measures yields the vector space of signed measures that are generated by $\mathcal{P}$, denoted $\mathcal{M}$. We assume that the topological vector space $\mathcal{M}$ is endowed with a topology that satisfies the following two properties: (1) the mappings

$$\begin{array}{rcl}(p, p') & \to & p + p' \\ (\alpha, p) & \to & \alpha p\end{array}$$

are continuous for every pair of priors $p, p' \in \mathcal{P}$ and scalar $\alpha \in \mathbb{R}$; and (2) a subset $A \subseteq \mathbb{R}$ is Borel if and only if for every pair of priors $p, p' \in \mathcal{P}$ the one-dimensional set of weighted averages

$$\{\alpha p + (1 - \alpha) p' : \alpha \in A\} \tag{2}$$

is a Borel subset of $\mathcal{M}$. These two properties are satisfied for a large variety of topologies on $\mathcal{M}$, including the topology of weak convergence and the topology of the total variation norm, but not for extremely strong topologies such as the totally disconnected topology in which every subset of $\mathcal{M}$ is open. The result below applies to any metric topology on $\mathcal{M}$ which satisfies the two properties above and that is also at least as strong as the topology of weak convergence.[11]

We start with two lemmata.

Lemma 1. Let $f_1, f_2 \in \mathcal{B}$ be two different BDP priors for bidder i, and let $f_1 - f_2 = f^+ - f^-$ be the Jordan decomposition of the signed measure $f_1 - f_2$ on T, where f^+ and f^- are two mutually singular positive measures on T. Then both f^+ and f^- are positive multiples of BDP priors for player i.

Proof. Because f_1 and f_2 satisfy (1), or

$$\int_{T_i} \left(\int_{T_{-i}} \varphi(\tau_i, \tilde{\tau}_{-i})\, d\widehat{b_i}(\tau_i)(\tilde{\tau}_{-i}) \right) df_{i|T_i}(\tau_i) = \int_T \varphi(\tau)\, df_i(\tau) \tag{3}$$

for every player $i \in N$ and every continuous $\varphi : T \to \mathbb{R}$, so does $f_1 - f_2$. Because $f_1 - f_2 = f^+ - f^-$ and both f^+, f^- are mutually singular positive measures, both f^+, f^- satisfy (3) as well. Therefore, $\frac{f^+}{|f^+|}, \frac{f^-}{|f^-|} \in \Delta(T)$ are common priors.

We next show that $\frac{f^+}{|f^+|}, \frac{f^-}{|f^-|}$ are BDP priors for player i. Because f_k, $k = 1, 2$ are BDP priors for player i, there exist subsets $T^{f_k} \subset T$ such that

$$f_k\left(T^{f_k}\right) = 1,$$

[11] In particular, if a topology of strategic proximity as discussed in the introduction belongs to this range of topologies, then our result implies that the shy subset of FSE priors cannot be open in that topology. Hence, only the set of NFSE priors remains a potential candidate for being open and dense in such a topology.

the projection $T_i^{f_k}$ of T^{f_k} on T_i is the graph of a function

$$\Phi_i^{f_k} : B_i^{f_k} \to S_i$$

(where $B_i^{f_k}$ is the projection of $T_i^{f_k}$ on $\Delta(T_{-i})$), and for every $\tau = (\ldots, \tau_j, \ldots) \in T^{f_k}$ and every player $j \neq i$, the marginal $\widehat{b}_j(\tau_j)_{|T_i}$ of τ_j on T_i assigns probability 1 to $T_i^{f_k}$. This means that if

$$T^{f_1} \cap T^{f_2} \neq \emptyset$$

then for every $\tau = (\ldots, \tau_j, \ldots) \in T^{f_1} \cap T^{f_2}$ and every player $j \neq i$, the marginal $\widehat{b}_j(\tau_j)_{|T_i}$ of τ_j on T_i assigns probability 1 to $T_i^{f_1} \cap T_i^{f_2}$. It then follows that the graphs of $\Phi_i^{f_1}$ and $\Phi_i^{f_2}$ coincide on $B_i^{f_1} \cap B_i^{f_2}$ almost surely according to both $f^1_{|T_i}$ and $f^2_{|T_i}$, because on $T^{f_1} \cap T^{f_2}$ each of the priors is an average of j's beliefs $\left\{\widehat{b}_j(\tau_j) : \tau = (\ldots, \tau_j, \ldots) \in T^{f_1} \cap T^{f_2}\right\}$. So if we define $\Phi_i : B_i^{f_1} \cup B_i^{f_2} \to S_i$ by

$$\Phi_i(b_i) = \begin{cases} \Phi_i^{f_1}(b_i) & b_i \in B_i^{f_1} \\ \Phi_i^{f_2}(b_i) & \text{otherwise} \end{cases}$$

then both $f^1_{|T_i}$ and $f^2_{|T_i}$ assign probability 1 to the graph of Φ_i.

This implies that the marginal of the signed measure $f_1 - f_2$ on T_i assigns measure zero to the complement of the graph of Φ_i. So this must also be true for the marginals of the finite, positive measures f^+ and f^- on $B_i \times S_i$. In other words, both the marginals of $\frac{f^+}{|f^+|}$ and $\frac{f^-}{|f^-|}$ on T_i assign probability 1 to the graph of Φ_i, which means that $\frac{f^+}{|f^+|}$ and $\frac{f^-}{|f^-|}$ are BDP priors. ■

Lemma 2. The set $\mathcal{B}$ of BDP priors for bidder i is a Borel subset of the space of priors $\mathcal{P}$.

Proof. If the lemma obtains when $\mathcal{P}$ is equipped with the topology of weak convergence, it also obtains for any stronger metric topology. It is therefore enough to proceed assuming that $\mathcal{P}$ is equipped with the topology of weak convergence.

By definition, a prior $p \in \mathcal{P}$ is a BDP prior if and only if the marginal of p on $T_i = V_i \times \Delta(T_{-i})$ is concentrated on a measurable graph $\Phi_i^p : B_i^p \to V_i$. This is expressible by countably many conditions, in the following way.

Since V_i is separable, there is a countable collection $\{A_i^n\}_{n\geq 1}$ of subsets of V_i which is closed under complements and finite unions and generates the Borel sigma-field of V_i. Hence there are also countably many partitions $\{\Gamma_i^m\}_{m\geq 1}$ of V_i to finitely many disjoint subsets $\left\{A_i^{n_k^m}\right\}_{k=1}^{N_i^m} \subseteq \{A_i^n\}_{n\geq 1}$. Similarly, Since $\Delta(T_{-i})$ is separable, there exists a countable collection $\left\{Y_i^\ell\right\}_{\ell\geq 1}$ of subsets of $\Delta(T_{-i})$ which is closed under complements and finite unions and generates the Borel sigma-field of $\Delta(T_{-i})$. Hence, there are also countably many partitions $\{\Lambda_i^r\}_{r\geq 1}$ of $\Delta(T_{-i})$ to finitely many disjoint subsets in $\left\{Y_i^{\ell_k^r}\right\}_{k=1}^{L_i^r} \subseteq \left\{Y_i^\ell\right\}_{\ell\geq 1}$.

The marginal of p on $T_i = V_i \times \Delta(T_{-i})$ is concentrated on the graph of Φ_i^p if and only if for every partition $\Gamma_i^m = \left\{A_i^{n_k^m}\right\}_{k=1}^{N_i^m}$ of S_i

$$p\left(\bigcup_{k=1}^{N_i^m}\left(A_i^{n_k^m} \times \left(\Phi_i^p\right)^{-1}\left(A_i^{n_k^m}\right) \times T_{-i}\right)\right) = 1$$

Intuitively, as the partitions $(\Gamma_i^m)_{m\geq 1}$ of V_i get finer, the union of the rectangles $A_i^{n_k^m} \times \left(\Phi_i^p\right)^{-1}\left(A_i^{n_k^m}\right)$ approximates increasingly well the graph of Φ_i^p.

Now, for each partition $\Gamma_i^m = \left\{A_i^{n_k^m}\right\}_{k=1}^{N_i^m}$ of V_i, $\left\{\left(\Phi_i^p\right)^{-1}\left(A_i^{n_k^m}\right)\right\}_{k=1}^{N_i^m}$ is a partition of $\Delta(T_{-i})$, that can be approximated arbitrarily well (in terms of the probabilities assigned to the partition members by the marginal of p on $\Delta(T_{-i})$) by partitions in $\{\Lambda_i^r\}_{r\geq 1}$. Hence, the marginal of p on $T_i = V_i \times \Delta(T_{-i})$ is concentrated on a measurable graph from $\Delta(T_{-i})$ to V_i if and only if for every natural number $q \geq 1$ and for each partition $\Gamma_i^m = \left\{A_i^{n_k^m}\right\}_{k=1}^{N_i^m}$ of V_i there exists a partition $\Lambda_i^r = \left\{Y_i^{\ell_k^r}\right\}_{k=1}^{L_i^r}$ of $\Delta(T_{-i})$ with $L_i^r = N_i^m$ and

$$p\left(\bigcup_{k=1}^{N_i^m}\left(A_i^{n_k^m} \times Y_i^{\ell_k^r} \times T_{-i}\right)\right) \geq 1 - \frac{1}{q}$$

Formally, therefore, the set $\mathcal{F}$ of FSE priors is

$$\bigcap_{i\in N}\bigcap_{m\geq 1}\bigcap_{q\geq 1}\bigcup_{r\geq 1}\left\{p \in \mathcal{P} : p\left(\bigcup_{k=1}^{N_i^m}\left(A_i^{n_k^m} \times Y_i^{\ell_k^r} \times T_{-i}\right)\right) \geq 1 - \frac{1}{q}\right\}$$

which is a Borel subset of the space of priors $\mathcal{P}$. ■

Theorem 1. The set $\mathcal{B}$ of BDP priors for bidder i is finitely shy in the space $\mathcal{P}$ of priors on the universal type space.

Proof. Let $g \in \mathcal{P}$ be a non-BDP prior for player i, and let $c \in \mathcal{P}$ such that c, g are mutually singular. Consider the one-dimensional subspace of $\mathcal{M}$

$$H = \{\alpha(g - c) : \alpha \in \mathbb{R}\}.$$

By (2), Lebesgue measure λ_H is well defined on H. We have that $\alpha(g-c) + c = \alpha g + (1-\alpha)c \in \mathcal{P}$ if and only if $\alpha \in [0,1]$ and hence $\lambda_H(\mathcal{P} - c) = 1 > 0$. Moreover, $H \setminus \{0\}$ consists entirely of signed measures $\alpha(g-c)$, $\alpha \neq 0$ whose marginals on T_i are not concentrated on a graph of some measurable function $\Phi : \Delta(T_{-i}) \to S_i$, because the marginal of g is not concentrated on such a graph by assumption, and $g \perp c$.

However, $\lambda_H(\mathcal{B}+x)=0$ for every $x\in\mathcal{M}$. In fact, $H\cap(\mathcal{B}+x)$ is either empty or a singleton. Indeed, assume by contradiction that

$$\begin{aligned} f_1+x &= h_1=\alpha_1(g-c) \\ f_2+x &= h_2=\alpha_2(g-c) \end{aligned}$$

where $h_1,h_2\in H$, $f_1,f_2\in\mathcal{B}$ and $\alpha_1>\alpha_2$. Then

$$f_1-f_2=(\alpha_1-\alpha_2)\,g-(\alpha_1-\alpha_2)\,c$$

By lemma 1, in the Jordan decomposition

$$f_1-f_2=f^+-f^-$$

f^+,f^- are non-negative multiples of BDE priors for player i. However, since both $(\alpha_1-\alpha_2)\,g$ and $(\alpha_1-\alpha_2)\,c$ are mutually singular positive measures, by the uniqueness of the Jordan decomposition we must have

$$\begin{aligned} f^+ &= (\alpha_1-\alpha_2)\,g \\ f^- &= (\alpha_1-\alpha_2)\,c \end{aligned}$$

But this is impossible, since $(\alpha_1-\alpha_2)\,g$ is a positive multiple of an non-BDE prior for player i. ■

Corollary 1. The set $\mathcal{F}$ of Full-Surplus-Extraction priors for player i is finitely shy in the set $\mathcal{P}$ of all common priors on the universal space T.

Proof. By proposition 1 we have $\mathcal{F}\subseteq\mathcal{B}$, so the corollary follows from the definition of finitely shy sets. ■

Remark 3. Inspection of the argument presented in this section reveals that it also implies that the set $\mathcal{F}_f$ of finite-support FSE priors for player i is finitely shy in the space of all finite-support private values priors $\mathcal{P}_f$ on the universal space T. The proofs apply verbatim.

Finally, we have shown in this section that the set of priors on the universal type space that permits the extraction of the players' full surplus is shy. The question of whether or not it is generically possible to approximate full surplus remains open. We conjecture, but have been so far unable to prove, that for every $\varepsilon>0$, both the sets of priors in which it is possible to extract up to ε of total surplus, and no more than ε below total surplus are large in the sense that neither is shy.[12]

[12] A partial answer to this question is provided by Neeman (2001), who for the case of public good provision, describes an example in which, if belief do not determine preferences, then the probability that the public good can be provided decreases to zero with the number of players while efficiency requires that the public good be provided with probability 1. This result may be interpreted as implying that the total surplus that can be extracted from the players converges to zero at the same time that the total surplus that could be generated by the players remains uniformly bounded away from zero.

3 Implementation with Interdependent Valuations

In this section, we demonstrate how the results obtained in the previous section for an auction problem with private values, can be generalized to any mechanism design problem with interdependent types. Specifically, we ask whether a given decision rule, which is a mapping from players' types into outcomes is generically implementable. Whenever possible, we rely on the notation used in the previous section.

Let $N = \{1, ..., n\}$ be a finite set of players, and $\mathcal{X}$ a measurable set of outcomes. The players' preferences over outcomes depend on the state of nature $k \in K$. The space of states of nature K is our *basic space of uncertainty.* It is assumed to be a complete, separable, metric space, that is endowed with its Borel σ-field. When the state of nature is $k \in K$, the outcome $x \in \mathcal{X}$ prevails, and player i receives a monetary transfer m_i, her payoff is given by

$$u_i(x, k) + m_i$$

where

$$u_i : \mathcal{X} \times K \to \mathbb{R}$$

is a Borel measurable function. The players are assumed to be expected utility maximizers.

3.1 Type Spaces

For every player $i \in N$, the set of player i's types Θ_i is assumed to be a complete, separable, metric space. Every type $\theta_i \in \Theta_i$ is associated with a probability measure on the space of states of nature K and the other players' types $\Theta_{-i} = \prod_{j \neq i} \Theta_j$. The space of probability measures $\Delta(K \times \Theta_{-i})$ is endowed with the topology of weak convergence. With a slight abuse of notation, we say that $\Theta_i \subseteq \Delta(K \times \Theta_{-i})$.

This formulation, which implies that the uncertainty of $\theta_i \in \Theta_i$ is about $K \times \Theta_{-i}$ but not about Θ_i, captures the idea that each type has a sufficiently developed introspective ability to determine its own belief.

Type θ_i's *belief type* $\widehat{b}_i(\theta_i) \in \Delta(\Theta_{-i})$ is the marginal $\theta_{i|\Theta_{-i}}$ of the probability measure θ_i on the other players' types Θ_{-i}. Type θ_i's *preference type* $\widehat{v}_i(\theta_i)$ is any version of the expected payoff functions

$$U_i\left(\tilde{x}; \theta_i, \tilde{\theta}_{-i}\right) : \mathcal{X} \times \Theta_{-i} \to \mathbb{R}, \quad \theta_i \in \Theta_i$$

that satisfies

$$\int_{\Theta_{-i}} U_i\left(\delta\left(\theta_i, \tilde{\theta}_{-i}\right); \theta_i, \tilde{\theta}_{-i}\right) d\theta_{i|\Theta_{-i}} = \int_{K \times \Theta_{-i}} u_i\left(\delta\left(\theta_i, \tilde{\theta}_{-i}\right), \tilde{\kappa}\right) d\theta_i$$

for every measurable *decision rule* $\delta : \Theta \to \mathcal{X}$.

The type space is the product $\Theta \equiv \prod_{i \in N} \Theta_i$ of the players' type sets. Each $\theta \in \Theta$ is called a state of the world.

The private-values setting of the previous section is a particular case of the formulation described in this section. If a type θ_i is a product probability

$$\theta_i = \theta_{i|K} \times \theta_{i|\Theta_{-i}}$$

then the expected payoff functions $U_i\left(\tilde{x};\theta_i,\tilde{\theta}_{-i}\right)$ are independent of $\tilde{\theta}_{-i}$, and can therefore be denoted $U_i\left(\tilde{x};\theta_i\right)$.

In the setting of the single object private values auction considered in the previous section, outcomes $x = (x_1, ..., x_n)$ are given by vectors that describe the probability with which each bidder or player wins the object. Bidders' payoffs are linear in the probability with which they win the object and are independent of both other bidders' types and the probabilities with which other bidders win the object. Hence, if we let $\mathbf{1}_i$ denote the vector that has 1 in the i-th place and 0 everywhere else, then for every vector x, $U_i\left(x;\theta_i\right) = x_i \cdot U_i\left(\mathbf{1}_i;\theta_i\right)$. Moreover, because $U_i\left(\mathbf{1}_i;\theta_i\right)$ describes θ_i's payoff when it wins the object for sure, which we denoted in the previous section by $\widehat{v}_i\left(\theta_i\right)$, every type's preferences can be completely described by its preference type $\widehat{v}_i\left(\theta_i\right)$.

To further illustrate the definition, consider now the case of a single object pure common value auction with two bidders. In this case $k \in K$ is the true value of the object. Suppose that bidder 1 knows the value k with certainty; that bidder 2 has no private information, only a belief which specifies the probabilities p_k of the potential values of k; and that all of this is common knowledge among the bidders. Then bidder 2 has a single type $\bar{\theta}_2$. Bidder 1's types have the form[13]

$$\theta_1^k = \delta_k \times \delta_{\bar{\theta}_2}, \quad k \in K,$$

respectively. Bidder 1's preference and belief types are therefore given by

$$\begin{aligned}
\widehat{v}_1\left(\theta_1^k\right) &= U_1\left(x;\theta_1^k,\bar{\theta}_2\right) = x_1 k \\
\widehat{b}_1\left(\theta_1^k\right) &= \delta_{\bar{\theta}_2}
\end{aligned}$$

The unique type $\bar{\theta}_2$ of bidder 2 is a probability measure over $K \times \Theta_1$, which assigns probability p_k to the combination $\left(k,\theta_1^k\right)$ for $k \in K$, and zero probability to any other combination in $K \times \Theta_1$. The preference and belief types of $\bar{\theta}_2$ are given by

$$\begin{aligned}
\widehat{v}_2\left(\bar{\theta}_2\right) &= U_2\left(x;\bar{\theta}_2,\theta_1^k\right) = x_2 k \\
\widehat{b}_2\left(\bar{\theta}_2\right)\left(\theta_1^k\right) &= p_k
\end{aligned}$$

Thus bidder 2's preference type depends non-trivially on bidder 1's type θ_1^k.

3.2 The Universal Type Space

Given the basic space of uncertainty K and the set of players N, there exists a *universal type space*

$$T = \prod_{i \in N} T_i$$

[13] δ denotes the unit-mass probability measure.

into which every other type space T can be uniquely mapped in a beliefs-preserving way (Mertens and Zamir, 1985; Brandenburger and Dekel, 1993; and Heifetz, 1993). That is, for every type space Θ there exists a unique set of measurable mappings[14] $(\mathbb{E}_i : \Theta_i \to T_i)_{i \in N}$ satisfying

$$\mathbb{E}_i(\theta_i)(A) = \theta_i\left(\mathbb{E}_{-i}^{-1}(A)\right)$$

for every measurable $A \subseteq K \times T_{-i}$, where

$$\mathbb{E}_{-i} : K \times \Theta_{-i} \to K \times T_{-i}$$

is defined by

$$\mathbb{E}_{-i}\left(k, (\theta_j)_{j \neq i}\right) = \left(k, (\mathbb{E}_j(\theta_j))_{j \neq i}\right)$$

It turns out that in the universal type space for player i, T_i, is isomorphic with $\Delta(K \times T_{-i})$ (and not just with a subset of it) for every $i \in N$. We therefore refer to T_i and $\Delta(K \times T_{-i})$ interchangeably.

The private-values universal type space T^{PV} described in the previous section is a subset of the universal type space T that is presented here in the special case in which $K = \prod_{i \in N} V_i$. It is the subset of T in which it is commonly known that each player i's types $\tau_i \in T_i$, $i \in N$, have a product form as follows

$$\tau_i = \tau_{i|V_i} \times \tau_{i|V_{-i} \times T_{-i}}.$$

The definition of a common prior on the universal space is the same as in section 2.3 above. The space of environments of interest is the set of priors $\mathcal{P}$ on the universal type space T.

3.3 EDR Priors are Non-Generic

In section 2 we have considered private-values environments of a particular kind, in which the private preferences of a player could be represented by a one-dimensional valuation. We now proceed to the general (quasi-linear) setup defined at the beginning of this section.

Definition. A prior p permits the implementation of a decision rule $\delta : T \to \mathcal{X}$ if there exists an incentive compatible and individually rational direct revelation mechanism $\left\langle \delta, (m_i)_{i \in N} \right\rangle$ where $m_i : T \to \mathbb{R}$ denotes the payment to player i as a function of the players' types, that implements δ.[15] A prior that permits the implementation of every decision rule is called an

[14] which are in fact also continuous.

[15] A direct revelation mechanism $\left\langle \delta, (m_i)_{i \in N} \right\rangle$ is incentive compatible if

$$\int_{T_{-i}} (U_i(\delta(\tau_i, \tilde{\tau}_{-i})) + m_i(\tau_i, \tilde{\tau}_{-i}))\, d\hat{b}_i(\tau_i)(\tilde{\tau}_{-i}) \geq \int_{T_{-i}} (U_i(\delta(\tau_i', \tilde{\tau}_{-i})) + m_i(\tau_i', \tilde{\tau}_{-i}))\, d\hat{b}_i(\tau_i)(\tilde{\tau}_{-i})$$

for every player $i \in N$, and player i's types τ_i, $\tau_i' \in T_i$. It is individually-rational if

$$\int_{T_{-i}} (U_i(\delta(\tau_i, \tilde{\tau}_{-i})) + m_i(\tau_i, \tilde{\tau}_{-i}))\, d\hat{b}_i(\tau_i)(\tilde{\tau}_{-i}) \geq 0$$

for every player $i \in N$, and player i's types $\tau_i \in T_i$.

EDR prior.

Definition. A prior p permits the full extraction of the players surplus relative to a decision rule $\delta : T \to \mathcal{X}$ if there exists an incentive compatible and individually rational direct revelation mechanism $\left\langle \delta, (m_i)_{i \in N} \right\rangle$ that implements δ with payment functions that leave each type of each player with zero surplus.[16]

Remark 4. As in the previous section, the revelation principle implies that in the two definitions above, no loss of generality is entailed by restricting attention to direct revelation mechanisms.

Remark 5. The previous section was devoted to investigating the possibility of the extraction of the players' full surplus relative to the ex-post efficient allocation rule in the context of a single object private values auction.

Remark 6. A prior p that permits full extraction relative to a decision rule δ also permits the implementation of δ but the opposite need not be true. For example, the second price auction implements the ex-post efficient allocation rule in a single object auction environment with private values, but as we have seen in the previous section, it is not generically the case that it is possible to extract the full surplus of the bidders relative to the ex-post efficient allocation rule. Hence the statement that it is impossible to implement a given decision rule is stronger than the statement that it is impossible to extract the players' surplus relative to this rule.

Recently, Aoyagi (1998) and d'Aspremont, Crémer, and Gérard-Varet (2002) showed that in models with at least 3 players and a fixed finite number of types for each player that is larger than or equal to 2, it is generically possible to implement every decision rule.[17] In contrast, we show below that finite-support EDR priors are BDP priors. Since the set $\mathcal{B}_f$ of finite-support BDP priors is finitely shy within the set $\mathcal{P}_f$ of *all* finite-support priors, we conclude that the set $\mathcal{E}_f$ of finite-support EDR priors is non-generic within $\mathcal{P}_f$.

In order to establish this result, we impose the mild assumption that for every player i there exists an outcome x_0^i that if implemented, generates a payoff of 0 for player i regardless of player i's type. Letting the players "opt out" of the mechanism ensures the existence of such outcomes.

Proposition 2. A finite support EDR prior is a BDP prior.

[16] That is, the payment functions $(m_i)_{i \in N}$ are such that

$$\int_{T_{-i}} m_i\left(\tau_i, \tilde{\tau}_{-i}\right) d\hat{b}_i\left(\tau_i\right)\left(\tilde{\tau}_{-i}\right) = -\int_{T_{-i}} U_i\left(\delta\left(\tau_i, \tilde{\tau}_{-i}\right)\right) d\hat{b}_i\left(\tau_i\right)\left(\tilde{\tau}_{-i}\right)$$

for every player $i \in N$, and player i's types $\tau_i \in T_i$.

[17] Their notion of implementation does not require individual rationality but rather budget-balance. However, every budget balanced mechanism can be transformed into an individualy rational mechanism by adding a sufficiently large constant to each player's payment function m_i.

Proof. Suppose that p is not a BDP prior for player i. Then player i has two types $\tau_i, \tau_i' \in T_i$ that each have a positive p-probability, the same beliefs

$$\hat{b}_i(\tau_i) = \hat{b}_i(\tau_i') \equiv b_i \in \Delta(T_{-i}),$$

but different preference types

$$\hat{v}_i(\alpha_i') \neq \hat{v}_i(\alpha_i).$$

That is, there exist a profile of other players' types $\overline{\tau}_{-i} \in T_{-i}$ such that

$$b_i(\overline{\tau}_{-i}) > 0$$

and an outcome $\bar{x} \in \mathcal{X}$ such that (without loss of generality)

$$U_i(\bar{x}; \tau_i, \overline{\tau}_{-i}) > U_i(\bar{x}; \tau_i', \overline{\tau}_{-i}) \geq 0.$$

Define the decision rule

$$\delta : T \to \mathcal{X}$$

by

$$\delta(\tau) = \begin{cases} \bar{x} & \tau = (\tau_i', \overline{\tau}_{-i}) \\ x_0^i & \text{otherwise} \end{cases}$$

Consider any system of monetary transfers

$$m_i : T \to \mathbb{R}, \quad i \in N$$

and suppose that the mechanism $\left\langle \delta, (m_i)_{i \in N} \right\rangle$ is incentive compatible. In particular, for type τ_i'

$$\sum_{\tilde{\tau}_{-i} \in T_{-i}} \left(U_i\left(\delta\left(\tau_i', \tilde{\tau}_{-i}\right); \tau_i', \tilde{\tau}_{-i}\right) + m_i\left(\tau_i', \tilde{\tau}_{-i}\right)\right) b_i\left(\tilde{\tau}_{-i}\right)$$
$$\geq \sum_{\tilde{\tau}_{-i} \in T_{-i}} \left(U_i\left(\delta\left(\tau_i, \tilde{\tau}_{-i}\right); \tau_i', \tilde{\tau}_{-i}\right) + m_i\left(\tau_i, \tilde{\tau}_{-i}\right)\right) b_i\left(\tilde{\tau}_{-i}\right)$$

or

$$U_i\left(\bar{x}; \tau_i', \tilde{\tau}_{-i}\right) b_i\left(\tilde{\tau}_{-i}\right) + \sum_{\tilde{\tau}_{-i} \in T_{-i}} m_i\left(\tau_i', \tilde{\tau}_{-i}\right) b_i\left(\tilde{\tau}_{-i}\right) \geq \sum_{\tilde{\tau}_{-i} \in T_{-i}} m_i\left(\tau_i, \tilde{\tau}_{-i}\right) b_i\left(\tilde{\tau}_{-i}\right).$$

Therefore, if instead of truthfully reporting its type, τ_i reports it is type τ_i', then its expected payoff is

$$\begin{aligned}
& \sum_{\tilde{\tau}_{-i} \in T_{-i}} \left(U_i\left(\delta\left(\tau_i', \tilde{\tau}_{-i}\right); \tau_i, \tilde{\tau}_{-i}\right) + m_i\left(\tau_i', \tilde{\tau}_{-i}\right)\right) b_i\left(\tilde{\tau}_{-i}\right) \\
= \; & U_i\left(\bar{x}; \tau_i, \overline{\tau}_{-i}\right) b_i\left(\overline{\tau}_{-i}\right) + \sum_{\tilde{\tau}_{-i} \in T_{-i}} m_i\left(\tau_i', \tilde{\tau}_{-i}\right) b_i\left(\tilde{\tau}_{-i}\right) \\
> \; & U_i\left(\bar{x}; \tau_i', \overline{\tau}_{-i}\right) b_i\left(\overline{\tau}_{-i}\right) + \sum_{\tilde{\tau}_{-i} \in T_{-i}} m_i\left(\tau_i', \tilde{\tau}_{-i}\right) b_i\left(\tilde{\tau}_{-i}\right) \\
\geq \; & \sum_{\tilde{\tau}_{-i} \in T_{-i}} m_i\left(\tau_i, \tilde{\tau}_{-i}\right) b_i\left(\tilde{\tau}_{-i}\right) \\
= \; & \sum_{\tilde{\tau}_{-i} \in T_{-i}} \left(U_i\left(\delta\left(\tau_i, \tilde{\tau}_{-i}\right); \tau_i, \tilde{\tau}_{-i}\right) + m_i\left(\tau_i, \tilde{\tau}_{-i}\right)\right) b_i\left(\tilde{\tau}_{-i}\right)
\end{aligned}$$

in contradiction to the presumed incentive compatibility of $\left\langle \delta, (m_i)_{i\in N} \right\rangle$. It follows that the decision rule δ cannot be implemented, so p is not an EDR prior. ■

Corollary 2. The set of finite support EDR priors is finitely shy in the set of priors with finite support $\mathcal{P}_f$.

3.4 Implementation of Efficient Decision Rules

As explained in the previous subsection, if players' beliefs determine their preferences, then not only can a seller extract the full surplus, but any social choice function can be implemented. The fact that, as we have shown in the previous subsection, players' beliefs do not generically determine their preferences, implies that not all social choice functions can be implemented. This still leaves open the question of whether or not a specific decision rule can be generically implemented. An answer to this question can be provided for the case of efficient decision rules.

Implementation of a decision rule requires that the mechanism designer be able to induce players to reveal both the preference and belief components of their types. A player's beliefs about other players' types can always be fully extracted at a cost by standard arguments (see e.g. d'Aspremont and Gerart-Varet, 1979). Thus a mechanism designer may generally face a trade-off between the cost and benefit of extracting a player's belief. However, if the decision rule to be implemented is efficient, then as shown by Bergemann and Morris (2003), the players' beliefs can be extracted at no cost. It therefore follows that it is possible to provide a precise characterization of whether a given social choice function is (interim) implementable on arbitrary finite type spaces in terms of conditions that arise when looking at standard implementation in an environment with stochastically independent types (for details, see Bergemann and Morris, 2003). The gist of this characterization is that after a player's belief type has been costlessly extracted, then if the player's beliefs do not determine its preferences, the player has to be given some rent in order to induce it to reveal its payoff type truthfully, in a similar way to the rent players have to be given when their types are stochastically independent.

4 Related Literature

Following an example in Myerson (1981) that showed that a seller in an auction may be able to exploit the presence of correlation among bidders to extract the bidders' full surplus, Crémer and McLean (1985, 1988) showed that a monopolistic seller can generically extract the full surplus of risk neutral consumers and bidders, respectively, in models with a fixed number of types.[18] McAfee and Reny (1992) constructed a similar auction to the one described by Crémer and McLean that (approximately) extracts the full surplus of the bidders when the number of bidders' types is uncountably large, but did not explicitly address the

[18] Crémer and McLean (1988, Appendix B) have indicated how some of their results can be generalized to allow for a continuum of types.

issue of genericity. McAfee et al. (1989), Johnson et al. (1992), and Brusco (1998) have established related results in more specific contexts. For a general formulation of this result, which allows for a continuum of multidimensional, mutually payoff relevant, agents' types, see Johnson et al. (2002). Recently, Aoyagi (1998) and d'Aspremont et al. (2002) have used a similar argument to the one used by Crémer and McLean (1985, 1988) to show that it is generically possible to implement any decision rule in models with finite type spaces.[19]

A number of authors have argued that the conditions that are imposed in order to obtain these full-rent-extraction results, while standard in many applications, are nevertheless very strong. Crémer and McLean (1988) suggested that full rent extraction is not robust to the introduction of risk aversion or limited liability constraints, and emphasized the dependence of these results on the common prior assumption. Following their suggestion, Robert (1991) showed that for any given auction mechanism, when agents are risk averse or face limited liability constraints, the function that relates the common prior to the seller's profit and to total surplus (and hence also to the sum of information rents captured by the agents) is continuous in the prior. Since it is known that agents do obtain positive information rents in independent environments, Robert concluded that full information rent extraction also fails in "nearly independent" environments with risk averse agents or agents that face limited liability constraints. More recently, Laffont and Martimort (2000) have established the continuity of the mechanism's outcome function also for environments with risk-neutral agents who are not constrained by limited liability, but who may form collusive coalitions. Intuitively, the reason that full rent extraction fails under these circumstances is that the auction mechanisms that extract the full buyers' rent rely on lotteries whose variance increases to infinity at independence. Thus, in nearly independent environments, mechanisms that rely on such lotteries violate the buyers' limited liability or participation constraints. Because these lotteries also prescribe payments to and from agents that strongly depend on the actions of other agents, mechanisms that rely on such lotteries are highly susceptible to collusion among the agents, and fail in nearly independent environments where these payments are large.

Finally, this paper makes a contribution to the growing literature about robust mechanism design that has stemmed out of Robert Wilson's view that further progress in game theory depends "on succesive reduction in the base of common knowledge required to conduct useful analyses of practical problems" (Wilson, 1987).[20] As shown by Neeman (2001), full-surplus extraction hinges on the fact that it is commonly believed that a player's belief determines, or predicts with certainty, the player's preferences. Once this assumption is relaxed, the full surplus of the players cannot be extracted. This paper presents a model in which it is shown that it is generically incorrect to assume that the designer and players maintain such common belief assumptions.

[19] What turns out to be the necessary and sufficient condition for implementation of every decision rule was first introduced by d'Aspremont and Gérard-Varet (1982). d'Aspremont et al. (2002) demonstrate that this condition (condition B) is strictly weaker than Aoyagi's (1998) strict regularity condition which have been shown to be sufficient for implementation of every decision rule.

[20] See, e.g., Bergemann and Morris (2003), Chung and Ely (2004), Neeman (2001,2003), Weinstein and Yildiz (2004) and the references therein.

4.1 The Relationship to Crémer and McLean (1988)

Crémer and McLean (1988) showed that within the set of models with a *fixed* finite number of types $n_i \geq 2$ for each player i (or equivalently, within the set of priors that are supported on a *fixed* finite number of types $n_i \geq 2$ for each player i), the set of priors that permit full-surplus extraction from any bidder is generic. How come we get the opposite result when we consider the set of priors that are supported on *all* finite numbers of types, or with arbitrary finite supports?

The main reason is that the set of priors that are supported on a fixed finite number of types is not closed under averaging. For example, the average of the common priors that are represented by the two matrices

	$\tau_2 = (v_2, b_2)$	$\tilde{\tau}_2 = \left(\tilde{v}_2, \tilde{b}_2\right)$
$\tau_1 = (v_1, b_1)$	a	b
$\tilde{\tau}_1 = \left(\tilde{v}_1, \tilde{b}_1\right)$	c	d

	$\tau_2' = (v_2, b_2')$	$\tilde{\tau}_2' = \left(\tilde{v}_2, \tilde{b}_2'\right)$
$\tau_1' = (v_1, b_1')$	a'	b'
$\tilde{\tau}_1' = \left(\tilde{v}_1, \tilde{b}_1'\right)$	c'	d'

(where $a+b+c+d = a'+b'+c'+d' = 1$) is not the prior that is represented by the matrix

	$\tau_2'' = (v_2, b_2'')$	$\tilde{\tau}_2'' = \left(\tilde{v}_2, \tilde{b}_2''\right)$
$\tau_1'' = (v_1, b_1'')$	$\frac{1}{2}(a+a')$	$\frac{1}{2}(b+b')$
$\tilde{\tau}_1'' = \left(\tilde{v}_1, \tilde{b}_1''\right)$	$\frac{1}{2}(c+c')$	$\frac{1}{2}(c+c')$

but rather the following prior

	$\tau_2 = (v_2, b_2)$	$\tilde{\tau}_2 = \left(\tilde{v}_2, \tilde{b}_2\right)$	$\tau_2' = (v_2, b_2')$	$\tilde{\tau}_2' = \left(\tilde{v}_2, \tilde{b}_2'\right)$
$\tau_1 = (v_1, b_1)$	$\frac{1}{2}a$	$\frac{1}{2}b$	0	0
$\tilde{\tau}_1 = \left(\tilde{v}_1, \tilde{b}_1\right)$	$\frac{1}{2}c$	$\frac{1}{2}d$	0	0
$\tau_1' = (v_1, b_1')$	0	0	$\frac{1}{2}a'$	$\frac{1}{2}b'$
$\tilde{\tau}_1' = \left(\tilde{v}_1, \tilde{b}_1'\right)$	0	0	$\frac{1}{2}c'$	$\frac{1}{2}d'$

which is a prior that is supported on 8 rather than 4 states.

In particular, the average of the common priors that are represented by the two matrices

	$v_2 = 0$	$\tilde{v}_2 = 1$
$v_1 = 0$	$\frac{1}{2}$	0
$\tilde{v}_1 = 1$	0	$\frac{1}{2}$

	$v_2 = 0$	$\tilde{v}_2 = 1$
$v_1 = 0$	0	$\frac{1}{2}$
$\tilde{v}_1 = 1$	$\frac{1}{2}$	0

is not

	$v_2 = 0$	$\tilde{v}_2 = 1$
$v_1 = 0$	$\frac{1}{4}$	$\frac{1}{4}$
$\tilde{v}_1 = 1$	$\frac{1}{4}$	$\frac{1}{4}$

but rather

	$v_2 = 0$	$\tilde{v}_2 = 1$	$v_2 = 0$	$\tilde{v}_2 = 1$
$v_1 = 0$	$\frac{1}{4}$	0	0	0
$\tilde{v}_1 = 1$	0	$\frac{1}{4}$	0	0
$v_1 = 0$	0	0	0	$\frac{1}{4}$
$\tilde{v}_1 = 1$	0	0	$\frac{1}{4}$	0

(Notably, the latter average preserves the fact that it is common knowledge that bidders know each other's type, while the former does not.)

The situation here is similar to the one in the following example. Let $[0, 1]$ be a set of goods. Agents consume bundles consisting of finitely many goods, and they consider close-by goods to be close substitutes. Formally, a bundle is a positive measure with a finite support on $[0, 1]$ (the measure of each good indicating its quantity in the bundle), where this set of measures is equipped with the topology of weak convergence (so that, for instance, the sequence of bundles with one unit of the good '$\frac{1}{2}$' and one unit of the good '$\frac{1}{2} + \frac{1}{n}$' converges to the bundle with *two* units of the good '$\frac{1}{2}$').

The goods in the Cantor set $C \subset [0, 1]$ are 'radioactive', and bundles that contain radioactive goods are 'dangerous.' Other bundles are 'safe.'

Consider, first, the bundles with exactly k goods. These bundles can equivalently be represented as the subset of vectors $((g_1, q_1), \ldots (g_k, q_k)) \in [0, 1]^k \times \mathbb{R}^k_{++}$ indicating the k good-quantity pairs, where all the goods are distinct. This set of bundles B_k is an open, dense and full-Lebesgue-measure subset of $[0, 1]^k \times \mathbb{R}^k_{++}$. Within B_k, the set S_k of safe bundles is an open, dense and full-Lebesgue-measure subset. Notice, however, that unions of k-good bundles are typically not k-good bundles.

Next, consider the space B of *all* bundles with finitely many goods each. This space *is* closed under the operation of taking unions of bundles.

The set S of safe bundles is dense in B, because $[0, 1] \setminus C$ is open and dense, so one can approximate arbitrarily well a dangerous bundle with a safe bundle, by replacing each radioactive good with the same quantity of an arbitrarily close non-radioactive good. However, the set of dangerous bundles is also dense in B, because adding an arbitrarily small quantity of some radioactive good to a bundle makes the new bundle dangerous.

The set B is an infinite dimensional cone. The set S of safe bundles is the intersection of B with a *proper subspace* of the vector space $V(B)$ generated by B. In this sense, S is "small" within B.

If B were the cone of a finite-dimensional space, then a subset $S \subset B$ with these properties would have zero measure according to the Lebesgue measure on B. The notion of finite shyness captures the same idea when B is infinite dimensional. Finite shyness is implied by the fact that the co-dimension of S within B is positive. In fact, in this example the co-dimension of S in B is infinite.

The intuition regarding the 'smallness' of the safe bundles would be further strengthened if we were to allow also for infinite bundles, represented by any positive measure on $[0, 1]$. In this general setting, the set of dangerous measure-bundles assigning some positive weight to the Cantor set C (which has zero Lebesgue measure but the cardinality of the continuum) is indeed prevalent.

Similarly, the co-dimension of the finite-support full-surplus-extraction priors $\mathcal{F}_f$ within the set of all finite-support priors $\mathcal{P}_f$ is infinite. There is a continuum of *extreme* common priors (i.e., priors that are not convex combinations of other priors), that permit full surplus extraction, and a continuum of extreme common priors that do not permit full surplus extraction and in which some player earns a positive information rent.

These two complementary sets of extreme priors are both dense within the set of extreme priors, which means that the surplus that the seller can extract is highly discontinuous in the prior. If bidders' beliefs, which are represented by convex combinations of extreme priors, change even slightly, then the surplus extracted by the seller can change discontinuously in the change of the bidders' beliefs.

In 'quantitative' terms, however, if one holds the view that any prior could 'just as well' be represented by any convex combination of extreme priors, then it would be a-typical if absolutely none of the continuum of the dense extreme priors which do not permit full-rent-expropriation were not to get even the tiniest weight. Moreover, extending the notion of convex combination to allow for every prior that can be represented by a probability measures π over the set of *extreme* priors with *any* support, only bolsters the intuition that FSE priors are rare. Indeed, to get a FSE prior, the (possibly diffused) π would have to 'miss' the entire continuum of dense extreme non-FSE priors. This is captured by Corollary 1, which states that the set of FSE priors is shy.

References

Anderson, R. M. and W. R. Zame (2001) "Genericity with Infinitely Many Parameters," *Advances in Theoretical Economics* **1**, pp. 1-62.

Aoyagi, M. (1998) "Correlated types and Bayesian incentive compatible mechanisms with budget balance," *Journal of Economic Theory* 79, 142-151.

d'Aspremont, C., and L.-A. Gérard-Varet (1979) "Incentives and Incomplete Information," *Journal of Public Economics* 11, 25-45.

d'Aspremont, C., and L.-A. Gérard-Varet (1982) "Bayesian incentive compatible beliefs," *Journal of Mathematical Economics* 10, 83-103.

d'Aspremont, C., J. Crémer, and L.-A. Gérard-Varet (2002) "Balanced Bayesian mechanisms," *Journal of Economic Theory*, forthcoming.

Aumann, R. (1998) "Common priors: a reply to Gul," *Econometrica* **66**, 929-938.

Bergemann, D. and S. Morris (2003) "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.

Brusco, S. (1998) "Unique implementation of the full surplus extraction outcome in auctions with correlated types," *Journal of Economic Theory*, **80**, 185-200.

Christensen, J. P. R. (1974) *Topology and Borel Structure*, North Holland Mathematical Studies Vol. 10, Amsterdam: North-Holland.

Chung, K. and J. C. Ely (2004) "Foundations of dominant strategy mechanisms," mimeo, Boston and Northwestern Universities.

Crémer, J. and R. McLean (1985) "Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent," *Econometrica* **53**, 345-361.

Crémer, J. and R. McLean (1988) "Full extraction of the surplus in Bayesian and dominant strategy auctions," *Econometrica* **56**, 1247-1257.

Gul, F. (1998) "A comment on Aumann's Bayesian view," *Econometrica* **66**, 923-927.

Harsanyi, J. (1967, 1968) "Games with incomplete information played by Bayesian players," *Management Science* **14**, 159-182, 320-334, 486-502.

Heifetz, A. (1993) "The Bayesian Formulation of Incomplete Information – the Non-Compact Case," *International Journal of Game Theory* **21**, 329-338.

Hunt, B. R., T. Sauer and J. A. Yorke (1992) "Prevalence: A Translation-Invariant 'Almost Every' on Infinite-Dimensional Spaces," *Bulletin (New Series) of the American Mathematical Society* **27**, pp. 217-238.

Johnson, S., J. Pratt, and R. Zeckhauser (1992) "Efficiency Despite Mutually Payoff-Relevant Private Information: The Finite Case," *Econometrica* **58**, 873-900.

Johnson, S., N. Miller, J. Pratt, and R. Zeckhauser, Efficient Design with Interdependent Valuations and an Informed Center," mimeo, Kennedy School of Government, Harvard University, 2002

Kajii, A. and S. Morris (1998) "Payoff Continuity in Incomplete Information Games," *Journal of Economic Theory* **82**, 267-276.

Laffont, J.-J., and D. Martimort (2000) "Mechanism design with collusion and correlation," *Econometrica* **68**, 309-342.

McAfee, P. R. and P. J. Reny (1992) "Correlated information and mechanism design," *Econometrica* **60**, 395-421.

McAfee, P. R., J. McMillan, and P. J. Reny (1989) "Extracting the surplus in the common-value auction," *Econometrica* **57**, 1451-1459.

Mertens, J.-F. and S. Zamir (1985) "Formulation of Bayesian analysis for games with incomplete information," *International Journal of Game Theory* **10**, 619-632.

Morris, S. (1995) "The common prior assumption in economic theory," *Economics and Philosophy* 11, 227-253.

Morris, S. (2001) "Typical Types," mimeo, Yale University.

Myerson, R. (1981) "Optimal Auction Design, *Mathematics of Operations Research* **6**, 58-73.

Neeman, Z. (2001) "The Relevance of Private Information in Mechanism Design," *Journal of Economic Theory,* forthcoming.

Neeman, Z. (2003) "The Effectiveness of the English Auction," *Games and Economic Behavior* 43 (2003), 214-238.

Robert, J. (1991) "Continuity in auction design," *Journal of Economic Theory* **55**, 169-179.

Weinstein, J. and M. Yildiz (2004) "Finite-Order Implications of Any Equilibrium, mimeo, MIT.

Wilson, R. (1987) "Game-Theoretic Analyses of Trading Processes," in *Advances in Economic Theory: Fifth World Congress*, Ed. T. Bewley, Cambridge: Cambridge University Press.

Yamasaki, Y. (1985) *Measures on Infinite Dimensional Spaces,* World Scientific, Singapore.