

Genericity and Robustness of Full Surplus Extraction*

Yi-Chun Chen[†] Siyang Xiong[‡]

May 14, 2012

Abstract

We study whether priors that admit full surplus extraction (FSE) are generic, which becomes a gauge to evaluate the validity of the current mechanism design paradigm. We consider the space of priors on the universal type space, and thereby relax the assumption of a fixed finite number of types in [Cr mer and McLean \(1988\)](#). We show that FSE priors are topologically generic, contrary to the result of [Heifetz and Neeman \(2006\)](#) that FSE is generically impossible, both geometrically and measure-theoretically. Instead of using the BDP approach or convex combinations of priors adopted in [Heifetz and Neeman \(2006\)](#), we prove our genericity results by showing a robustness property of *Cr mer-McLean mechanisms*.

*We thank the editor and anonymous referees for insightful comments which greatly improve the paper. We thank Eddie Dekel, Jeff Ely, Eduardo Faingold, Songying Fang, Aviad Heifetz, Martin Hellwig, Jinwoo Kim, Maciej Kotowski, Takashi Kunimoto, Qingmin Liu, Xiao Luo, Daisuke Oyama, Wolfgang Pesendorfer, Marciano Siniscalchi, Yeneng Sun, Olivier Tercieux, and participants of seminars/conferences at Kansas Workshop on Economic Theory, Rice University, University of Tokyo, and 11th SAET Conference for helpful comments and discussions. All remaining errors are our own.

[†]Department of Economics, National University of Singapore, ecsycc@nus.edu.sg

[‡]Department of Economics, Rice University, xiong@rice.edu

1 Introduction

In economic models, agents with private information about their independent values retain some informational rent (Myerson (1981)). However, the source of the informational rent is not privacy; rather, it is the independence of information among agents. In a seminal paper, Crémer and McLean (1988) (hereafter, CM) show that in a classical mechanism design model (hereafter, the classical model), a mechanism designer can *fully* extract agents' rent even if their values are only slightly correlated. Since "nearly all" models have correlated information, full surplus extraction (FSE) should be a generic phenomenon. The result implies that private information is (generically) irrelevant, which seems patently false in practice. As McAfee and Reny wrote,

“The results (full rent extraction) cast doubt on the value of the current mechanism design paradigm as a model of institution design” (McAfee and Reny, 1992, p.400).¹

Many explanations have been offered to address this *FSE critique*. We know now that CM's genericity result does not hold if any of the following essential assumptions of the classical model is violated — risk neutrality, unlimited liability, absence of collusion among agents, and lack of competition among sellers (see Robert (1991), Laffont and Martimort (2000), and Peters (2001)). However, the classical model that incorporates *all* these assumptions is still commonly used. By modifying these essential assumptions of the classical model, the aforementioned explanations only demonstrate that the *modified* models are immune to the FSE critique, but do not explain why the classical model itself generates predictions that contradict our observations. To resolve the issue, we must be able to attribute the genericity of FSE to assumptions which are *inessential* to the classical model and yet are critical to the genericity of FSE. Failing to find such inessential assumptions would invite re-examination of the classical model and all the theories based on it.

¹CM prove their result in the single-object private-value auction setting with a fixed finite type space, but their result easily extends to general mechanism design problems. McAfee and Reny (1992) characterize FSE in a general mechanism design setting with a continuum of types whose beliefs are given by continuous density functions.

In an important paper, [Heifetz and Neeman \(2006\)](#) (hereafter, HN) identify such an assumption. In particular, HN point out that CM's genericity result hinges on their implicit common-knowledge assumption that each agent has a fixed finite number of types. However, there is no *a priori* finite bound on the number of types in mechanism design problems. This common-knowledge assumption is therefore inessential to the classical model. Relaxing this assumption, we should study the genericity/non-genericity of FSE in the space of general priors supported on an arbitrary number of types. Following this view, HN prove that FSE priors are "negligible" (non-generic) in a geometric sense (i.e., they are contained in a proper face), and a measure-theoretical sense (i.e., they are contained in a finitely shy set as defined in [Anderson and Zame \(2001\)](#)).²

In this paper, we also relax CM's common-knowledge assumption of a fixed finite number of types, and yet we prove that FSE is *topologically* generic. Namely, while the fixed finite number of types is an inessential assumption, it has no effect on the genericity of FSE in the topological sense. More importantly, our results imply that the classical model remains subject to the FSE critique.

We study private information modeled by (common) priors on the universal type space (see [Mertens and Zamir \(1985\)](#) and [Heifetz and Neeman \(2006\)](#)). The universal type space is a (Harsanyi) type space which embeds all possible type spaces. Therefore, our approach not only relaxes the assumption of a fixed finite number of types, but also entails no loss of generality (see Section 4.4). Following [Mertens, Sorin, and Zamir \(1994\)](#), we endow the space of priors with the standard weak* topology.³ We say a set is (topologically) generic if it is a residual set, i.e., it contains a countable intersection of open and

²Finite shyness is a notion which extends the notion of shyness which is originally proposed by [Christensen \(1974\)](#) and rediscovered by [Hunt, Sauer, and Yorke \(1992\)](#). See [Anderson and Zame \(2001\)](#) for details.

³Agents choose their best strategy according to their expected utility, and the mechanism designer chooses the optimal mechanism according to her expected revenue. The weak* topology is the coarsest topology which makes these "expected values" continuous in beliefs and it is often regarded as a coarse topology. [Barelli \(2009\)](#) and [Chen and Xiong \(2011\)](#) also adopt the weak* topology on priors. Moreover, the definition of the measure-theoretic notion of genericity adopted in HN also requires a topology on priors and HN take the weak* topology as an important candidate. See Section 4.3 for more discussion on other topologies.

dense sets; a set is non-generic if its complement is generic, i.e., it is a meager set.⁴

We report two genericity results. Say a prior is finite if its support is a finite set, and infinite otherwise. We first consider the space of all priors which can be finite or infinite. To address the technical difficulties associated with infinite priors, we follow [McAfee and Reny \(1992\)](#) (hereafter, MR) by considering the notion of *almost FSE* (AFSE). Say a prior is AFSE if for any $\varepsilon > 0$, on the support of the prior we can construct a mechanism which is Bayesian incentive compatible, and individually rational, and extracts $1 - \varepsilon$ proportion of the total surplus. Our first result shows that AFSE priors form a generic set in the space of all priors.⁵ We then restrict attention to the space of finite priors. Say a prior is FSE if on the support of the prior we can construct a mechanism which is dominant strategy incentive compatible, and individually rational, and extracts the total surplus. Our second result shows that FSE priors form a generic in the space of finite priors.

We also study an important subclass of priors called models. A model is a prior which is not a convex combination of other priors. We discuss a sense in which every mechanism design problem associated with a prior can be reduced to a mechanism design problem associated with models. We show that HN's geometric analysis and measure-theoretic analysis do not apply to the space of models, whereas our topological analysis does and our genericity results remain the same.

Unlike HN, we prove our genericity results by directly studying the mechanism design problems on general priors.⁶ We consider a class of mechanisms called *Crémer-McLean mechanisms* (CM mechanisms). A CM mechanism is a second-price auction sup-

⁴Being residual and meager are standard notions of topological genericity (being typical) and nongenericity (being negligible), respectively. This notion of genericity is also adopted in [Barelli \(2009\)](#), [Chen and Xiong \(2011\)](#), [Dekel, Fudenberg, and Morris \(2006\)](#), and [Ely and Pęski \(2011\)](#).

⁵That is, generically, AFSE can be partially implemented. In Section 4.2, we further show that, generically, AFSE can be fully virtually Bayesian implemented in the sense of [Abreu and Matsushima \(1992\)](#) and [Duggan \(1997\)](#).

⁶HN's genericity result relies upon a property of a prior called *BDP* which is due to [Neeman \(2004\)](#). A prior satisfies the BDP property if it assigns probability one to a set of type profiles in which no distinct types have the same belief. We show in [Chen and Xiong \(2011\)](#) that priors which satisfy the BDP property are topological generic. Since BDP is necessary but not sufficient for FSE or AFSE (see HN's Proposition 2), the genericity of BDP priors has no implication on the genericity or non-genericity of FSE or AFSE priors.

plemented with side payments (see Definition 9). The gist of our proof is a robustness property of CM mechanisms: For any $\varepsilon > 0$, if a CM mechanism extracts $1 - \varepsilon$ proportion of the total surplus for a prior μ , it would continue to do so for *any* prior in a small weak*-neighborhood of μ . The notable feature of this robustness property is that it applies to any priors in a *large* space (i.e., the universal space) which approximates μ in a *weak* notion of proximity (i.e., the weak* topology).

Given this robustness property, the intuition of our genericity results becomes transparent. For any $\varepsilon > 0$, let \mathcal{A}_ε denote the set of priors for which more than $1 - \varepsilon$ proportion of the total surplus can be extracted via CM mechanisms. \mathcal{A}_ε is open by the robustness property; \mathcal{A}_ε is also dense because it contains all FSE priors which are known to be dense by the existing literature. That is, the set of priors for which FSE is approximately achieved (i.e., \mathcal{A}_ε for any arbitrarily small ε) is both open and dense. Since the set of AFSE priors contains the residual set $\bigcap_{n=1}^{\infty} \mathcal{A}_{1/n}$, it follows that AFSE is generic.

Economic modelers obviously would not want to rule out the correlated private information, which admits FSE/AFSE, as proved in CM and MR. In this paper, we strengthen their results by showing that the mechanisms they employ would do almost equally well for all nearby priors. Consequently, for a mechanism design problem in a classical model with a slightly misspecified prior, the solution under correlated information (in the sense of CM and MR) is robust, while the solution under independent information is not robust. This posts a fundamental restriction for economic modelers — revenue maximization aligns with the robustness only under correlated information.⁷

⁷The following quote illustrates how our results strengthen the FSE critique:

“It is a reasonable position that in the analysis of a social or physical system, the properties one should first focus on are those that enjoy persistency, that is, stability under perturbations, and are typical - informally, those whose qualitative characteristics do not depend too precisely on the environmental variables (persistency) and that hold but "exceptionally" in all admissible environments (typicality). The underlying justification for both desiderata is the same: In a world that is not observed, or perhaps not even given to us, in a very precise manner, only the persistent and typical have a good chance to be observed” (Mas-Colell, 1985, pp.316-317).

That is, our results show that according the classical model, the FSE property is the persistent (robust) and typical (generic) that "have a good chance to be observed."

Finally, we note that the opposite conclusions in HN and this paper are solely due to the different perspectives we take, i.e., HN take the geometric/measure-theoretical approach, whereas we take the topological approach. Since the space of priors is an infinite-dimensional space for which there is no consensus on the notion of genericity, which one is more appropriate depends on the specific context.⁸ Despite the difference, our result shows that the measure-theoretic non-genericity of FSE represents a knife-edge situation because for any $\varepsilon > 0$, the set \mathcal{A}_ε , being open and dense, is not negligible even in the measure-theoretic sense.⁹

The rest of this paper is organized as follows. For simplicity, we focus on the single-object private-value auction environment in our main results throughout the main text and we discuss general mechanism design problems in Section 4.1. Section 2 contains notations and definitions. Section 3 presents our main results. Section 4 discusses related issues. Section 5 concludes.

2 Preliminaries

Throughout this paper, for any compact metric space X with the metric d_X , we endow X with the Borel σ -algebra. Let $\Delta(X)$ denote the topological space of all probability measures on the Borel σ -algebra with the weak* topology. The weak* topology is metrizable under the Prohorov distance defined as

$$\rho(\mu, \mu') = \inf \{ \varepsilon > 0 : \mu(E) \leq \mu'(E^\varepsilon) + \varepsilon, \forall \text{Borel set } E \subset X \}, \forall \mu, \mu' \in \Delta(X). \quad (1)$$

where $E^\varepsilon \equiv \{x' : \inf_{x \in E} d_X(x', x) < \varepsilon\}$ (see (Dudley, 2002, 11.3.3. Theorem)). All product spaces are endowed with the product topology and subspaces with the relative topology. The support of a probability measure $\mu \in \Delta(X)$, denoted by $\text{Supp}\mu$, is the intersection of all closed sets with measure one under μ . For any finite set $F \subset X$, let $|F|$ denote the cardinality of F .

⁸Anderson and Zame (2001) point out some weakness of the residual (resp. meager) set as the notion of genericity (resp. nongenericity). Stinchcombe (2000) discusses some caveats of the prevalence (resp. shy) set as the notion of genericity (resp. nongenericity).

⁹This follows because an open set cannot be finitely shy (see Anderson and Zame (2001)). Thus, this observation holds in any topology which is finer than the weak* topology.

2.1 Priors and belief spaces

One object is for sale. Let I be a finite set of bidders. For simplicity, we assume that for every $i \in I$, $V_i = [0, 1]$ is the set of bidder i 's possible values of the object endowed with the Euclidean topology. Let Θ_i^* be the compact metric private-value universal type space on $V = \prod_{i \in I} V_i = [0, 1]^{|I|}$, which contains all possible bidder i 's types (see [Mertens and Zamir \(1985\)](#), ([Heifetz and Neeman, 2006](#), pp.228-229), and also Section 4.4). Let $v_i : \Theta_i^* \rightarrow V_i$ and $b_i : \Theta_i^* \rightarrow \Delta(\Theta_{-i}^*)$ be the continuous functions through which each $\theta_i \in \Theta_i^*$ identifies a value $v_i(\theta_i)$ and a belief $b_i(\theta_i)$ of bidder i . Let $\Theta^* = \prod_{i \in I} \Theta_i^*$ be the space of all bidders' type profiles and $\Theta_{-i}^* = \prod_{j \neq i} \Theta_j^*$ be the space of bidder i 's opponents' type profiles. Let d_i denote the metric on Θ_i^* , d_{-i} the metric on Θ_{-i}^* , and d the metric on Θ^* , where $d_{-i}(\theta_{-i}, \theta'_{-i}) = \max_{j \neq i} d_j(\theta_j, \theta'_j)$ and $d(\theta, \theta') = \max_{j \in I} d_j(\theta_j, \theta'_j)$. For each $\theta \in \Theta^*$, we denote by θ_i the type of bidder i under θ and we often save the notation to write $v_i(\theta)$ and $b_i(\theta)$ instead of $v_i(\theta_i)$ and $b_i(\theta_i)$.

A *belief subspace* Θ is a nonempty and compact subset of Θ^* such that for any $\theta \in \Theta$, $\{\theta_i\} \times \text{Supp} b_i(\theta)$ is a subset of Θ . A belief subspace is said to be *minimal* if it does not contain any belief subspace as a proper subset. For any $\mu \in \Delta(\Theta^*)$, denoted by μ_i the marginal distribution of μ on Θ_i^* . A probability measure $\mu \in \Delta(\Theta^*)$ is said to be a (common) *prior* if for any bounded measurable function $\varphi : \Theta^* \rightarrow \mathbb{R}$,

$$\int_{\Theta_i^*} \left(\int_{\Theta_{-i}^*} \varphi(\theta_i, \theta'_{-i}) b_i(\theta_i) [d\theta'_{-i}] \right) \mu_i [d\theta_i] = \int_{\Theta^*} \varphi(\theta) \mu [d\theta], \forall i. \quad (2)$$

Let $\mathcal{P} \subset \Delta(\Theta^*)$ be the set of all priors endowed with the weak* topology metrized by the Prohorov metric $d_{\mathcal{P}}$. The support of a prior μ is a belief subspace (see ([Mertens, Sorin, and Zamir, 1994](#), p.147, item 2)) and it is denoted by Θ^μ .¹⁰ A *finite prior* is a prior whose support is a finite set. A *model* is a prior whose support is a minimal belief subspace. Let \mathcal{P}^f denote the space of finite priors. Let \mathcal{M} denote the space of models and \mathcal{M}^f denote the space of finite models. For any set $E \subset \Theta^*$, let E_i and E_{-i} denote the projections of E into Θ_i^* and Θ_{-i}^* , respectively. The following result will be used later.

¹⁰In the literature, a belief subspace which is the support of some common prior is called a *consistent belief subspace*. Throughout the paper, we restrict our attention to consistent belief subspaces, and we omit "consistent" for simplicity.

Lemma 1 ((Mertens, Sorin, and Zamir, 1994, Theorem 3.1)) \mathcal{P}^f is dense in \mathcal{P} .

2.2 Mechanisms and surplus extraction

A mechanism designer tries to sell the object to the agents (i.e., bidders) in I . By the revelation principle, it is without loss of generality to restrict attention to direct mechanisms. A (direct) mechanism on a belief subspace Θ is a list of measurable functions $(q, m) \equiv (q_i : \Pi_{i \in I} \Theta_i \rightarrow [0, 1], m_i : \Pi_{i \in I} \Theta_i \rightarrow \mathbb{R})_{i \in I}$ satisfying $\sum_{i \in I} q_i(\theta) \leq 1$ for any $\theta \in \Pi_{i \in I} \Theta_i$. For each profile of reports θ , $q_i(\theta)$ specifies the probability that bidder i wins the object and $m_i(\theta)$ specifies how much bidder i pays.

Fix any mechanism (q, m) defined on a belief subspace Θ . Let $u_i(\theta'_i, \theta_{-i} | \theta_i, q, m)$ denote the (expected) payoff of type θ_i , when he reports θ'_i and the other bidders report $\theta_{-i} \in \Theta_{-i}$, i.e.,

$$u_i(\theta'_i, \theta_{-i} | \theta_i, q, m) \equiv v_i(\theta_i) q_i(\theta'_i, \theta_{-i}) - m_i(\theta'_i, \theta_{-i}). \quad (3)$$

Furthermore, let $U_i(\theta'_i | \theta_i, q, m)$ denote the interim expected payoff of type θ_i when he reports θ'_i and the other bidders truthfully reveal their types, i.e.,

$$U_i(\theta'_i | \theta_i, q, m) = \int_{\Theta_{-i}} u_i(\theta'_i, \theta_{-i} | \theta_i, q, m) b_i(\theta_{-i}) [d\theta_{-i}].$$

To simplify our notation, we write $U_i(\theta_i | q, m)$ for $U_i(\theta_i | \theta_i, q, m)$, which is the interim expected payoff of type θ_i when all bidders truthfully reveal their types.

Definition 1 A mechanism (q, m) defined on a belief subspace Θ is individually rational (IR) if

$$U_i(\theta_i | q, m) \geq 0 \text{ for every } (i, \theta) \in I \times \Theta.$$

Definition 2 A mechanism (q, m) defined on a belief subspace Θ is dominant strategy incentive compatible (DSIC) if

$$u_i(\theta_i, \theta_{-i} | \theta_i, q, m) \geq u_i(\theta'_i, \theta_{-i} | \theta_i, q, m) \text{ for every } (i, \theta, \theta') \in I \times \Theta \times \Theta. \quad (4)$$

That is, (q, m) on Θ satisfies DSIC iff truthful reporting is a weakly dominant strategy for every bidder.¹¹

Definition 3 A mechanism (q, m) defined on a belief subspace Θ is Bayesian incentive compatible (BIC) if

$$U_i(\theta_i | q, m) \geq U_i(\theta'_i | \theta_i, q, m) \text{ for every } (i, \theta, \theta') \in I \times \Theta \times \Theta.$$

That is, (q, m) on Θ satisfies BIC iff truthful reporting constitutes a Bayesian Nash equilibrium on Θ . Clearly, if (q, m) is DSIC, then it is also BIC.

Definition 4 For any prior μ and $\varepsilon \geq 0$, a mechanism (q, m) defined on Θ^μ achieves ε -surplus-extraction (ε -SE) if

$$\int_{\Theta^\mu} \left[\max_{i \in I} v_i(\theta) - \sum_{i \in I} m_i(\theta) \right] \mu[d\theta] \leq \varepsilon.$$

The maximal social surplus is $\int_{\Theta^\mu} \max_{i \in I} v_i(\theta) \mu[d\theta]$, while the surplus collected by the seller is $\int_{\Theta^\mu} \sum_{i \in I} m_i(\theta) \mu[d\theta]$. Under a mechanism that achieves ε -SE, at most ε surplus are surrendered to the bidders. We now define two notions of full surplus extraction.

Definition 5 A prior μ is a full-surplus-extraction (FSE) prior if there exists a mechanism (q, m) on Θ^μ which achieves IR, DSIC and 0-SE.

Definition 6 A prior μ is an almost-full-surplus-extraction (AFSE) prior if for any $\varepsilon > 0$, there exists a mechanism (q, m) on Θ^μ which achieves IR, BIC, and ε -SE.

We use \mathcal{F} and \mathcal{F}^f to denote the sets of FSE priors in \mathcal{P} and \mathcal{P}^f , respectively. Thus, $\mathcal{F}^f = \mathcal{F} \cap \mathcal{P}^f$. We use \mathcal{A} to denote the set of AFSE priors in \mathcal{P} . We require DSIC for FSE and BIC for AFSE. We make the difference so as to accommodate some technical difficulties associated with infinite priors.

¹¹(4) is usually called *ex post* Incentive Compatibility (EPIC). In the private-value environment considered in this paper, DSIC and EPIC are equivalent (see Bergemann and Morris (2005)).

Throughout the paper, we focus on a class of mechanisms defined as follows. Let Θ be a belief subspace. The *first-order belief* of $\theta_i \in \Theta_i$, denoted by $b_i^1(\theta_i) \in \Delta(V)$, is the belief of θ_i over the value profiles of all bidders. That is,

$$b_i^1(\theta_i) [\bar{V}] = b_i^1(\theta_i) [\{\theta_{-i} \in \Theta_{-i} : (v_i(\theta_i), v_{-i}(\theta_{-i})) \in \bar{V}\}], \forall \text{Borel set } \bar{V} \subset V.$$

Definition 7 A mechanism (q, m) defined on a belief subspace Θ is said to be a *first-order mechanism* if for every $\theta, \theta' \in \Theta$, $b_i^1(\theta_i) = b_i^1(\theta'_i)$ for all $i \in I$ implies $(q_i(\theta), m_i(\theta)) = (q_i(\theta'), m_i(\theta'))$ for all $i \in I$.

That is, the allocation and the payment of a first-order mechanism depend only on the first-order beliefs of the reported types. One mechanism that will play a crucial role in our analysis is the second-price auction (with an arbitrary tie-breaking rule) which we denote by (q^*, m^*) . That is, (q^*, m^*) is defined on Θ^* such that

$$\begin{aligned} \max_{i \in I} v_i(\theta) &= \sum_{i \in I} v_i(\theta_i) q_i^*(\theta_i, \theta_{-i}); \\ m_i^*(\theta) &= q_i^*(\theta) \times \max_{j \neq i} v_j(\theta). \end{aligned} \quad (5)$$

It is well known that (q^*, m^*) is IR and DSIC (and hence also BIC). Furthermore, (q^*, m^*) is a first-order mechanism. Indeed, (q^*, m^*) depends only on the reported values, i.e., for every $\theta, \theta' \in \Theta$, $v_i(\theta_i) = v_i(\theta'_i)$ for all $i \in I$ implies $(q_i^*(\theta), m_i^*(\theta)) = (q_i^*(\theta'), m_i^*(\theta'))$ for all $i \in I$. Since $b_i^1(\theta_i) = b_i^1(\theta'_i)$ implies $v_i(\theta_i) = v_i(\theta'_i)$, (q^*, m^*) is a first-order mechanism.

2.3 Genericity

In a topological space, a set is said to be nowhere dense if its closure has no interior point. A countable union of nowhere dense sets is called a meager set. The complement of a meager set is called a residual set. That is, a residual set contains a countable intersection of open and dense sets.¹²

¹²The notions of residual sets and meager sets are usually studied on Baire spaces. A Baire space is a topological space in which every nonempty open set is not meager. A residual set in a Baire space is dense and is not meager. The spaces of priors that we study in this paper are all Baire spaces.

Definition 8 In a topological space X , we say $F (\subset X)$ is generic in X if F is a residual set in X . We say $E (\subset X)$ is non-generic in X if E is a meager set in X .

The following technical result will be used later, and its proof is relegated to Appendix [A.2](#).

Lemma 2 If $Y (\subset X)$ is dense in X and $U (\subset X)$ is generic in X , then $U \cap Y$ is generic in Y under the relative topology.

3 Main results

In this section, we present our main results. We first define a class of mechanisms called Crémer-McLean mechanisms (CM mechanisms), which is a special class of first-order mechanisms. Then, we establish two robustness properties of CM mechanisms (Propositions [2-3](#)). We use these robustness properties to prove that almost full surplus extraction is generic in the space of all priors (Theorem [1](#)), and that full surplus extraction is generic in the space of finite priors (Theorem [2](#)). Finally, we prove these results for the space of models (Theorems [3-4](#)).

3.1 CM mechanisms

Recall that (q^*, m^*) denotes the second-price auction defined on Θ^* .

Definition 9 A mechanism (q, m) defined on Θ^* is called a Crémer-McLean (CM) mechanism if there exists a profile of continuous functions $(w_i : V_{-i} \rightarrow \mathbb{R})_{i \in I}$ such that $q_i(\theta) = q_i^*(\theta)$ and $m_i(\theta) = m_i^*(\theta) + w_i(v_{-i}(\theta_{-i}))$ for every $(i, \theta) \in I \times \Theta^*$.

The function w_i is often called a side-payment scheme for bidder i .^{[13](#)} It has three distinct features. First, w_i depends only on the reported values, which implies that a CM

¹³A CM mechanism is always defined on Θ^* . Alternatively, we may define a CM mechanism on a subspace Θ^μ . Or equivalently, the side payment may be defined only on $v_{-i}(\Theta_{-i}^\mu)$ rather than on V_{-i} . How-

mechanism depends only on the reported values. Consequently, a CM mechanism is a first-order mechanism. Second, w_i depends only on the reports of bidders $-i$ but not on bidder i 's report. As a result, truthful reporting remains a weakly dominant strategy for every type in a CM mechanism, as in a second-price auction. Finally, since (q^*, m^*) is fixed on Θ^* , a CM mechanism can be identified by the side-payment scheme. Hence, we write w for a CM mechanism and write $(w_i)_{i \in I}$ or simply w_i for side payments. The proof of the following lemma is straightforward and it is omitted.

Lemma 3 *Every CM mechanism is a first-order mechanism that satisfies DSIC. Furthermore, a prior μ admits FSE under a CM mechanism w if and only if $U_i(\theta_i|w) = 0$ for every $(i, \theta) \in I \times \Theta^\mu$.*

Define

$$\mathcal{F}^{cm} \equiv \{\mu \in \mathcal{P} : \mu \text{ admits FSE in a CM mechanism}\}.$$

Note that a CM mechanism may not satisfy IR. The full surplus extraction literature starting from [Cr mer and McLean \(1988\)](#) seeks to identify a side-payment scheme which extracts exactly the right amount of the surplus and maintains the IR constraint, i.e., to achieve $U_i(\theta_i|w) = 0$ for every $(i, \theta) \in I \times \Theta^\mu$. Define

$$\mathcal{P}_n^f \equiv \left\{ \mu \in \mathcal{P}^f : |\Theta_i^\mu| = n, \forall i \in I \right\}.$$

That is, \mathcal{P}_n^f is the set of priors whose supports are belief subspaces containing exactly n types for each bidder. Pick any $\mu \in \mathcal{P}_n^f$. Define the interim belief matrix for bidder i as

$$\mathbf{B}_i^\mu \equiv [b_i(\theta_i) [\theta_{-i}]]_{\theta_i \in \Theta_i^\mu, \theta_{-i} \in \Theta_{-i}^\mu},$$

where each row corresponds to a type θ_i of bidder i and her belief about the distribution of her opponents' types, i.e., $[b_i(\theta_i) [\theta_{-i}]]_{\theta_{-i} \in \Theta_{-i}^\mu}$. We say \mathbf{B}_i^μ has *full rank* if its column space has rank n , and μ has full rank if \mathbf{B}_i^μ has full rank for every $i \in I$. The FSE result in [Cr mer and McLean \(1988\)](#) implies the following lemma.

ever, our current definition incurs no loss of genericity: By Tietze Extension Theorem (see [Aliprantis and Border, 2006](#), 2.47)) and closedness of $v_{-i}(\Theta_{-i}^\mu)$, for any continuous $w'_i : v_{-i}(\Theta_{-i}^\mu) \rightarrow \mathbb{R}$, there exists a continuous $w_i : V_{-i} \rightarrow \mathbb{R}$ such that $w_i(\hat{v}_{-i}) = w'_i(\hat{v}_{-i})$ for any $\hat{v}_{-i} \in v_{-i}(\Theta_{-i}^\mu)$.

Lemma 4 $\mu \in \mathcal{F}$ if $\mu \in \mathcal{P}_n^f$ has full rank.

Clearly, the set of priors $\mathcal{P}_n^f \cap \mathcal{F}^f$ that have full rank is open and dense in \mathcal{P}_n^f . This is precisely the genericity result in [Cr mer and McLean \(1988\)](#). However, as [Heifetz and Neeman \(2006\)](#) argue, there is no *a priori* finite bound for the number of types when we model a situation involving asymmetric information. Hence, we should relax this common-knowledge assumption on the size of priors' supports, which leads us to study the genericity of FSE in \mathcal{P}^f and \mathcal{P} .

We introduce two more definitions. First, say that a finite prior μ has *full support* if $\Theta^\mu = \prod_{i \in I} \Theta_i^\mu$, i.e., the support of μ is a product set. Note that every finite prior that has full support is also a model. Second, say a finite prior μ has *distinct values* iff $v_i(\theta) = v_j(\theta')$ implies $i = j$ and $\theta_i = \theta'_j$. Define

$$\mathcal{F}^* \equiv \left\{ \mu \in \bigcup_{n=1}^{\infty} \mathcal{P}_n^f : \mu \text{ has full rank, full support, and distinct values} \right\}. \quad (6)$$

Thus, $\mathcal{F}^* \subset \mathcal{M}^f$. Now consider any $\mu \in \mathcal{F}^*$. Since μ has full rank, we can employ the result in [Cr mer and McLean \(1988\)](#) to construct a profile of functions $\tilde{w} = (\tilde{w}_i : \Theta_{-i}^\mu \rightarrow \mathbb{R})_{i \in I}$ which extracts all the surplus. Moreover, since μ has all distinct values, $v_{-i}(\theta_{-i}) \neq v_{-i}(\theta'_{-i})$ for every $\theta, \theta' \in \Theta^\mu$. Then, by Tietze Extension Theorem (see [Aliprantis and Border, 2006, 2.47](#)), there exists a CM mechanism $w = (w_i : V_{-i} \rightarrow \mathbb{R})_{i \in I}$ such that $w_i(v_{-i}(\theta)) = \tilde{w}_i(\theta)$, i.e., FSE is achieved under w . As a result $\mathcal{F}^* \subset \mathcal{F}^{cm}$. We thus have the following result.

Lemma 5 $\mathcal{F}^* \subset \mathcal{F}^{cm} \cap \mathcal{M}^f$.

The next result shows that \mathcal{F}^* is dense in \mathcal{P} . Note that [Lemma 5](#) and [Lemma 6](#) imply that $\mathcal{F}^{cm} \cap \mathcal{M}^f$ is dense in $\mathcal{M}^f, \mathcal{P}^f, \mathcal{M}$ and \mathcal{P} .

Lemma 6 \mathcal{F}^* is dense in \mathcal{P} .

Lemma 6 is an immediate consequence of the following facts: i) \mathcal{F}^* is dense in \mathcal{P}^f ,¹⁴ ii) \mathcal{P}^f is dense in \mathcal{P} (Lemma 1).

3.2 Robustness of CM mechanisms

For any mechanism (q, m) defined on Θ^* and $\varepsilon > 0$, define

$$\Omega_i(\varepsilon|q, m) \equiv \{\theta_i \in \Theta_i^* : U_i(\theta_i|q, m) \in (0, \varepsilon)\}, \forall i \in I.$$

That is, $\Omega_i(\varepsilon|q, m)$ is the set of i 's types whose interim expected payoffs under (q, m) fall between 0 and ε . To show the robustness property of CM mechanisms, we need the following lemmas.

Lemma 7 *Let w be a CM mechanism defined on Θ^* . Then,*

(a) *for every $i \in I$, $U_i(\theta_i|w)$ is a continuous function in θ_i ;*

(b) *for any $\varepsilon, \varepsilon' > 0$, $\{\mu \in \mathcal{P} : \mu_i[\Omega_i(\varepsilon|w)] > 1 - \varepsilon', \forall i \in I\}$ is open.*

Lemma 7 has two parts: part (a) says that the interim payoffs of the bidders under a CM mechanism are continuous, which follows immediately from the continuity of v_i , w_i and b_i ; part (b) says that if μ assigns high probability to types which retain only a small surplus, it is still the case for μ' sufficiently close to μ . Part (b) follows from part (a) and the definition of the Prohorov metric $d_{\mathcal{P}}$. The proof of Lemma 7 is standard and can be found in [Chen and Xiong \(2012\)](#).

Recall that a CM mechanism may not satisfy IR. Lemma 8 provides a way to modify a CM mechanism so as to achieve IR. The proof is relegated to Appendix A.1.

¹⁴To see the idea, consider the case with two bidders. In the following examples, $\mu \in \mathcal{P}^f$ and $\mu \notin \mathcal{F}^*$, while $\mu_m \in \mathcal{F}^*$ for all m large and $\mu_m \rightarrow \mu$ as $m \rightarrow \infty$.

	μ	$v_2 = 0$	$v_2 = 1/2$	$v_2 = 1$	
	$v_1 = \frac{1}{m}$	$\left(1 - \frac{1}{m}\right) \times \frac{1}{4}$	$\frac{1}{16m}$	$\frac{1}{16m}$	
$\mu :$	$v_1 = 0$	$\frac{1}{4}$	0	0	and $\mu_m :$
	$v_1 = 1$	0	$\frac{3}{8}$	$\frac{3}{8}$	
	$v_1 = 1 + \frac{1}{m}$	$\frac{1}{16m}$	$\left(1 - \frac{1}{m}\right) \times \frac{3}{16}$	$\left(1 + \frac{1}{m}\right) \times \frac{3}{16}$	
		$\frac{1}{16m}$	$\left(1 + \frac{1}{m}\right) \times \frac{3}{16}$	$\left(1 - \frac{1}{m}\right) \times \frac{3}{16}$	

Lemma 8 Let w be a CM mechanism defined on Θ^* and (q, m) be a mechanism defined on Θ^* as follows.

$$(q_i(\theta), m_i(\theta)) = \begin{cases} (q_i^*(\theta), m_i^*(\theta) + w_i(v_{-i}(\theta))), & \text{if } U_i(\theta_i|w) > 0; \\ (0, 0), & \text{if } U_i(\theta_i|w) \leq 0. \end{cases} \quad (7)$$

Then, (q, m) is a first-order mechanism which achieves IR and BIC on Θ^* .

We now state and prove two robustness results on FSE achieved in CM mechanisms. Both Proposition 1 and Proposition 2 say that if μ admits FSE in a CM mechanism, we can find a mechanism which extracts most of the surplus in a neighborhood around μ . The mechanism in Proposition 1 is a CM mechanism which is DSIC but need not be IR. In contrast, the mechanism in Proposition 2 respects IR but is only BIC.

Proposition 1 For any $\varepsilon, \varepsilon' > 0$ and any prior $\mu \in \mathcal{F}^{cm}$, there exist $\gamma > 0$ and a CM mechanism w such that for any μ' with $d_{\mathcal{P}}(\mu', \mu) < \gamma$, we have $\mu'_i[\Omega_i(\varepsilon|w)] > 1 - \varepsilon'$ for every $i \in I$.

Proof. Since $\mu \in \mathcal{F}^{cm}$, there is a CM mechanism w' such that $U_i(\theta_i|w') = 0$ for any $(i, \theta) \in I \times \Theta^\mu$. Define a new CM mechanism w such that

$$w_i(v_{-i}) = w'_i(v_{-i}) - \frac{\varepsilon}{4}, \forall (i, v) \in I \times V.$$

That is, the only difference between w and w' is that the side payment in w is always less than the side payment in w' by $\frac{\varepsilon}{4}$. Therefore, w is a CM mechanism such that $U_i(\theta_i|w) = \frac{\varepsilon}{4} \in (0, \frac{\varepsilon}{2})$ for any $(i, \theta) \in I \times \Theta^\mu$. Hence, $\mu_i[\Omega_i(\varepsilon|w)] = 1$ for every $i \in I$.

By Lemma 7(b), for any $\varepsilon' > 0$, $\{\mu' \in \mathcal{P} : \mu'_i[\Omega_i(\varepsilon|w)] > 1 - \varepsilon', \forall i \in I\}$ is open. Since $\mu \in \{\mu' \in \mathcal{P} : \mu'_i[\Omega_i(\varepsilon|w)] > 1 - \varepsilon', \forall i \in I\}$, there exists $\gamma > 0$ such that for any μ' with $d_{\mathcal{P}}(\mu', \mu) < \gamma$, we have $\mu'_i[\Omega_i(\varepsilon|w)] > 1 - \varepsilon'$ for every $i \in I$. ■

Proposition 2 For any $\varepsilon > 0$ and any prior $\mu \in \mathcal{F}^{cm}$, there exist $\gamma > 0$ and a first-order mechanism (q, m) defined on Θ^* such that for any μ' with $d_{\mathcal{P}}(\mu', \mu) < \gamma$, (q, m) achieves IR, BIC, and ε -SE on Θ^* .

Proof. By Proposition 1, there is a CM mechanism w defined on Θ^* and $\gamma > 0$ such that for any μ' with $d_{\mathcal{P}}(\mu', \mu) < \gamma$, we have

$$\mu'_i \left[\Omega_i \left(\frac{\varepsilon}{4|I|} |w \right) \right] > 1 - \frac{\varepsilon}{4|M||I|}, \quad \forall i \in I. \quad (8)$$

$$\text{where } M \equiv \max_{\theta_i \in \Theta_i^*} |U_i(\theta_i | w)| + 1. \quad (9)$$

By Lemma 8, we can construct a first-order mechanism (q, m) defined on Θ^* which achieves IR and BIC on Θ^* ; moreover, by (7),

$$U_i(\theta_i | q, m) = \max \{ U_i(\theta_i | w), 0 \}, \quad (10)$$

$$q_i(\theta) = q_i^*(\theta) \text{ if } \theta_i \in \Omega_i \left(\frac{\varepsilon}{4|I|} |w \right). \quad (11)$$

We are ready to show (q, m) achieves ε -SE on $\Theta^{\mu'}$ for any μ' with $d_{\mathcal{P}}(\mu', \mu) < \gamma$. First,

$$\begin{aligned} & \int_{\Theta^{\mu'}} \left[\sum_{i \in I} v_i(\theta_i) q_i(\theta_i, \theta_{-i}) - \sum_{i \in I} m_i(\theta) \right] \mu' [d\theta] \quad (12) \\ &= \sum_{i \in I} \int_{\Theta_i^{\mu'}} U_i(\theta_i | q, m) \mu'_i [d\theta_i] \\ &\leq \sum_{i \in I} \left[\int_{\Omega_i \left(\frac{\varepsilon}{4|I|} |w \right)} U_i(\theta_i | w) \mu'_i [d\theta_i] + \int_{\Theta_i^{\mu'} \setminus \Omega_i \left(\frac{\varepsilon}{4|I|} |w \right)} |U_i(\theta_i | w)| \mu'_i [d\theta_i] \right] \\ &\leq |I| \times \left[\frac{\varepsilon}{4|I|} + M \times \frac{\varepsilon}{4M|I|} \right] = \frac{\varepsilon}{2}, \end{aligned}$$

where the first inequality follows from (10); the second inequality follows from (8) and (9). Second,

$$\begin{aligned} & \int_{\Theta^{\mu'}} \left[\max_{i \in I} v_i(\theta) - \sum_{i \in I} v_i(\theta_i) q_i(\theta_i, \theta_{-i}) \right] \mu' [d\theta] \quad (13) \\ &= \int_{\Theta^{\mu'}} \left[\sum_{i \in I} v_i(\theta_i) q_i^*(\theta_i, \theta_{-i}) - \sum_{i \in I} v_i(\theta_i) q_i(\theta_i, \theta_{-i}) \right] \mu' [d\theta] \\ &= \int_{\Theta^{\mu'}} \left[\sum_{i \in I} v_i(\theta_i) [q_i^*(\theta_i, \theta_{-i}) - q_i(\theta_i, \theta_{-i})] \right] \mu' [d\theta] \\ &= \sum_{i \in I} \int_{\Theta_i^{\mu'} \setminus \Omega_i \left(\frac{\varepsilon}{4|I|} |w \right)} v_i(\theta_i) [q_i^*(\theta_i, \theta_{-i}) - q_i(\theta_i, \theta_{-i})] \mu'_i [d\theta_i] \\ &\leq \sum_{i \in I} \mu'_i \left[\Theta_i^{\mu'} \setminus \Omega_i \left(\frac{\varepsilon}{4|I|} |w \right) \right] \\ &\leq |I| \times \frac{\varepsilon}{4M|I|} \leq \frac{\varepsilon}{4}, \end{aligned}$$

where the first equality follows from (5); the third equality follows from (11); the first inequality follows from the fact that $v_i(\theta_i) [q_i(\theta_i, \theta_{-i}) - q_i^*(\theta_i, \theta_{-i})] \leq 1$ for all (i, θ) ; the second inequality follows from (8); the last equality follows because $M \geq 1$.

Combining (12) and (13), we get

$$\int_{\Theta^{\mu'}} \left[\max_{i \in I} v_i(\theta) - \sum_{i \in I} m_i(\theta) \right] \mu' [d\theta] \leq \varepsilon.$$

That is, (q, m) achieves ε -SE on $\Theta^{\mu'}$. ■

3.3 Genericity in the space of priors

3.3.1 Genericity of almost full surplus extraction in the space of priors

We first prove our genericity result in \mathcal{P} .

Theorem 1 *\mathcal{A} is generic in \mathcal{P} . That is, almost full surplus extraction is generically possible in the space of priors.*

Proof. Define

$$\mathcal{A}_n \equiv \left\{ \mu \in \mathcal{P} : \exists \text{ a first-order mechanism } (q, m) \text{ on } \Theta^\mu \text{ which achieves IR, BIC, and } \frac{1}{n}\text{-SE} \right\}. \quad (14)$$

Clearly, $\bigcap_{n=1}^{\infty} \mathcal{A}_n \subset \mathcal{A}$. Hence, it suffices to show that \mathcal{A}_n contains an open and dense set in \mathcal{P} .

Recall that $\mu \in \mathcal{F}^{cm}$ iff μ admits FSE in a CM mechanism. Then, by Proposition 2, for any $\mu \in \mathcal{F}^{cm}$, there exists $\gamma_\mu > 0$ and a first-order mechanism (q, m) defined on Θ^* such that for any μ' with $d_{\mathcal{P}}(\mu', \mu) < \gamma_\mu$, (q, m) achieves IR, BIC, and $\frac{1}{n}$ -SE on Θ^μ . Namely,

$$\mathcal{B}_{d_{\mathcal{P}}}(\mu, \gamma_\mu) \equiv \left\{ \mu' \in \mathcal{P} : d_{\mathcal{P}}(\mu', \mu) < \gamma_\mu \right\} \subset \mathcal{A}_n, \forall \mu \in \mathcal{F}^*.$$

Hence,

$$\mathcal{F}^{cm} \subset \bigcup_{\mu \in \mathcal{F}^{cm}} \mathcal{B}_{d_{\mathcal{P}}}(\mu, \gamma_\mu) \subset \mathcal{A}_n.$$

Since \mathcal{F}^{cm} is dense in \mathcal{P} by Lemmas 5-6, $\cup_{\mu \in \mathcal{F}^{cm}} \mathcal{B}_{d_{\mathcal{P}}}(\mu, \gamma_{\mu})$ is an open and dense set which is contained in \mathcal{A}_n . ■

3.3.2 Genericity of full surplus extraction in the space of finite priors

We next prove our genericity result in \mathcal{P}^f .¹⁵

Theorem 2 \mathcal{F}^f is generic in \mathcal{P}^f . That is, full surplus extraction is generically possible in the space of finite priors.

Proof. Define

$$\mathcal{A}_n^f = \left\{ \mu \in \mathcal{P}^f : \exists \text{ a CM mechanism } w \text{ on } \Theta^* \text{ s.t. } \mu_i \left[\Omega_i \left(\frac{1}{n} | w \right) \right] > 1 - \frac{1}{n}, \forall i \in I \right\}.$$

Step 1. \mathcal{A}_n^f contains an open and dense set in \mathcal{P}^f .

By Proposition 1, for each $\mu \in \mathcal{F}^{cm} \cap \mathcal{P}^f$, there exists $\gamma_{\mu} > 0$ and a CM mechanism w on Θ^* such that $\mu'_i \left[\Omega_i \left(\frac{1}{n} | w \right) \right] > 1 - \frac{1}{n}$ for any μ' with $d_{\mathcal{P}}(\mu', \mu) < \gamma_{\mu}$. Namely,

$$\mathcal{B}_{d_{\mathcal{P}}}(\mu, \gamma_{\mu}) \equiv \left\{ \mu \in \mathcal{P}^f : d_{\mathcal{P}}(\mu', \mu) < \gamma_{\mu} \right\} \subset \mathcal{A}_n^f, \forall \mu \in \mathcal{F}^{cm} \cap \mathcal{P}^f.$$

Hence,

$$\mathcal{F}^{cm} \cap \mathcal{P}^f \subset \cup_{\mu \in \mathcal{F}^{cm} \cap \mathcal{P}^f} \mathcal{B}_{d_{\mathcal{P}}}(\mu, \gamma_{\mu}) \subset \mathcal{A}_n^f.$$

Since $\mathcal{F}^{cm} \cap \mathcal{P}^f$ is dense in \mathcal{P}^f by Lemmas 5-6, $\cup_{\mu \in \mathcal{F}^{cm} \cap \mathcal{P}^f} \mathcal{B}_{d_{\mathcal{P}}}(\mu, \gamma_{\mu})$ is an open and dense set which is contained in \mathcal{A}_n^f .

Step 2. $\cap_{n=1}^{\infty} \mathcal{A}_n^f \subset \mathcal{F}^f$.

Fix any $\mu \in \cap_{n=1}^{\infty} \mathcal{A}_n^f$ and we will show $\mu \in \mathcal{F}^f$. Since μ is fixed in the proof, for simplicity, we write Θ for Θ^{μ} . Define

$$\begin{aligned} \mathbf{B}_i &\equiv [b_i(\theta_i) [\{\theta_{-i} \in \Theta_{-i} : v_{-i}(\theta_{-i}) = v_{-i}\}]]_{\theta_i \in \Theta_i, v_{-i} \in v_{-i}(\Theta_{-i})}; \\ \mathbf{U}_i^* &\equiv [U_i(\theta_i | q^*, m^*)]_{\theta_i \in \Theta_i}. \end{aligned}$$

¹⁵Theorem 2 and Lemma 2 imply that \mathcal{F} is not non-generic in \mathcal{P} . Whether \mathcal{F} is generic in \mathcal{P} remains an open question.

That is, \mathbf{B}_i is the belief matrix which summarizes bidder i 's beliefs about her opponents' values and \mathbf{U}_i^* is the column vector of the interim expected payoffs for bidders i ' types in the support of μ under the second-price auction (q^*, m^*) .

Since μ is a finite prior, Θ is a finite set. Define $\alpha \equiv \min_{(i,\theta) \in I \times \Theta} \mu_i[\{\theta_i\}] > 0$. For any n with $\frac{1}{n} < \alpha$, by the fact that $\mu \in \mathcal{A}_n^f$ and the definition of \mathcal{A}_n^f , there exists some CM mechanism w^n such that $\mu_i \left[\Omega_i \left(\frac{1}{n} | w^n \right) \right] > 1 - \frac{1}{n}$ for every $i \in I$. Since $\alpha = \min_{(i,\theta) \in I \times \Theta} \mu_i[\theta_i] > \frac{1}{n}$, it follows that $\mu_i \left[\Omega_i \left(\frac{1}{n} | w^n \right) \right] = 1$ for every $i \in I$, i.e., $U_i(\theta_i | w^n) \in \left(0, \frac{1}{n}\right)$ for every $(i, \theta) \in I \times \Theta$. That is,

$$\|\mathbf{U}_i^* - \mathbf{B}_i \mathbf{W}_i^n\|_\infty \in \left(0, \frac{1}{n}\right) \text{ for any } n > \frac{1}{\alpha}, \quad (15)$$

where \mathbf{W}_i^n is the column vector of side-payments $w_i^n(v_{-i})$ indexed by $v_{-i} \in v_{-i}(\Theta_{-i})$ and $\|\cdot\|_\infty$ is the supmetric in Euclidean space, i.e., $\|(a_1, \dots, a_k)\|_\infty = \max\{|a_1|, \dots, |a_k|\}$.

The vector $\mathbf{B}_i \mathbf{W}_i^n$ is contained in the column space of \mathbf{B}_i . Since (15) is true for any $n > \frac{1}{\alpha}$, \mathbf{U}_i^* is in the closure of the column space of \mathbf{B}_i . Since the column space of a finite matrix is a finite-dimensional linear space which is closed (see (Aliprantis and Border, 2006, 5.22 Corollary)), it follows that that \mathbf{U}_i^* is in the column space of \mathbf{B}_i . Namely, there exists $\bar{\mathbf{W}}_i = [\bar{w}_i(v_{-i})]_{v_{-i} \in v_{-i}(\Theta_{-i})}$ for any $i \in I$ such that $\mathbf{U}_i^* = \mathbf{B}_i \bar{\mathbf{W}}_i$. That is, $U_i(\theta_i | \bar{w}) = 0$ for every $(i, \theta) \in I \times \Theta$. Therefore, $\mu \in \mathcal{F}^f$ by Lemma 3. ■

3.4 Genericity in the space of models

We now show that our genericity results remain the same in the space of models.

Theorem 3 $\mathcal{A} \cap \mathcal{M}$ is generic in \mathcal{M} . That is, almost full surplus extraction is generically possible in the space of models.

Since \mathcal{M} is dense in \mathcal{P} by Lemma 5 and 6, Theorem 3 follows from Lemma 2 and Theorem 1.

Theorem 4 $\mathcal{F}^f \cap \mathcal{M}^f$ is generic in \mathcal{M}^f . That is, full surplus extraction is generically possible in the space of finite models.

Since \mathcal{M}^f is dense in \mathcal{P}^f by Lemma 5 and 6, Theorem 4 follows from Lemma 2 and Theorem 2.

The space of models is an important class of priors to which our topological analysis is applicable but HN's geometric and measure-theoretic analysis is not. The genericity notion employed by HN requires the ambient space of priors be convex. The space of all priors is indeed convex and in this case our topological genericity results stand in contrast to HN's non-genericity results. However, the space of models is not convex. For example, consider μ^a , which is a convex combination of two models μ' and μ'' described below.

$$\mu' : \begin{array}{|c|c|c|} \hline & \theta'_2 & \tilde{\theta}'_2 \\ \hline \theta'_1 & \frac{1}{4} & \frac{1}{4} \\ \hline \tilde{\theta}'_1 & \frac{1}{4} & \frac{1}{4} \\ \hline \end{array} \quad \mu'' : \begin{array}{|c|c|c|} \hline & \theta''_2 & \tilde{\theta}''_2 \\ \hline \theta''_1 & \frac{1}{6} & \frac{1}{3} \\ \hline \tilde{\theta}''_1 & \frac{1}{3} & \frac{1}{6} \\ \hline \end{array} \quad \mu^a : \begin{array}{|c|c|c|c|c|} \hline & \theta'_2 & \tilde{\theta}'_2 & \theta''_2 & \tilde{\theta}''_2 \\ \hline \theta'_1 & a \times \frac{1}{4} & a \times \frac{1}{4} & 0 & 0 \\ \hline \tilde{\theta}'_1 & a \times \frac{1}{4} & a \times \frac{1}{4} & 0 & 0 \\ \hline \theta''_1 & 0 & 0 & (1-a) \times \frac{1}{6} & (1-a) \times \frac{1}{3} \\ \hline \tilde{\theta}''_1 & 0 & 0 & (1-a) \times \frac{1}{3} & (1-a) \times \frac{1}{6} \\ \hline \end{array}$$

Observe that μ^a is not a model since its support contains two smaller belief subspaces.

In general, any prior can be regarded as a convex combination of models.¹⁶ HN interpret the convex combination as uncertainty faced by the mechanism designer. However, the designer can conceivably resolve this uncertainty and reduce a mechanism design problem associated with a prior to a mechanism design problem associated with models. To achieve this, the designer can simply ask the agents to simultaneously report the model containing their actual types. If their reports do not match, the designer levies a large fine. Since it is common knowledge among the agents that they are in the same model, truthfully reporting the true model is an equilibrium.¹⁷ For such a designer, the genericity issue is "how often" the resulting model permits FSE. Since the notions of genericity that HN adopt require convexity and we know of no other notion of genericity which does not require convexity, we can only apply the topological notion of genericity to the space of models.

¹⁶This follows from Choquet's theorem when we regard models as extreme points in the convex space of priors (see (Barelli, 2009, Section 4)). Using the observation, (Barelli, 2009, Section 4) also argue that HN's analysis is incomparable to CM's analysis, since the latter does not rely on this convex combination.

¹⁷See also (Chung and Ely, 2007, p.448) for a similar argument.

4 Discussion

4.1 General Mechanism Design Problems

Here we discuss how our results can be used to address surplus extraction in the general mechanism design problems formulated in [McAfee and Reny \(1992\)](#). Let $(X, (u_i)_{i \in I})$ be a general environment where X is the space of outcomes and $u_i : X \times V \times \mathbb{R} \rightarrow \mathbb{R}$ is a quasilinear ex post utility of agent i . Consider an arbitrary BIC mechanism (q', m') defined on Θ^* that gives agent i an interim expected equilibrium payoff of $U_i(\theta_i | q', m')$ (which we denote by $U'_i(\theta_i)$ for simplicity) when his type is $\theta_i \in \Theta_i^*$. Following [McAfee and Reny \(1992\)](#), we assume that U'_i is continuous.

For any continuous side payment scheme $(w_i)_{i \in I}$ on Θ^* , define

$$U'_i(\theta_i | w) \equiv U'_i(\theta_i) - \int w_i(v_{-i}(\theta_{-i})) b_i(\theta_i) [d\theta_{-i}].$$

Call the mechanism supplemented with the side-payment scheme $(w_i)_{i \in I}$ a CM' mechanism also denoted by w . For example, in our auction setting, the initial mechanism is the second-price auction and the resulting CM' mechanism is a CM mechanism, i.e., $U'_i(\theta_i) = U_i(\theta_i | q^*, m^*)$ and $U'_i(\theta_i | w) = U_i(\theta_i | w)$. Furthermore, any CM' mechanism is BIC.

In this setup, we define the notion of FSE' priors and AFSE' priors as follows.

- μ is an FSE' prior if there is a CM' mechanism w such that $U'_i(\theta_i | w) = 0$ for every $(i, \theta) \in I \times \Theta^\mu$;
- μ is an AFSE' prior if for any $\varepsilon > 0$, there is a CM' mechanism w on Θ^μ such that $\mu[\{\theta : U'_i(\theta_i | w) \in [0, \varepsilon], \forall i\}] \geq 1 - \varepsilon$.

Note that IR holds for an FSE' prior, while IR holds only with probability arbitrarily close to one for an AFSE' prior. In the following, we first explain that all of our results remain unchanged under these definitions, and then comment on the issue of IR for AFSE' priors. The first step does not require any restriction on the mechanism design environment studied in [McAfee and Reny \(1992\)](#), while the second step does.

First, we can replace $U_i(\theta_i|w)$ with $U'_i(\theta_i|w)$ and CM mechanisms with CM' mechanisms everywhere in Sections 3.2–3.3.2 and proceed as follows. Continuity of U'_i and w_i implies that $U'_i(\theta_i|w)$ is continuous and hence the thesis of Lemma 7. Proposition 1 then follows from Lemma 7. Now define

$$\mathcal{A}'_n \equiv \left\{ \begin{array}{l} \mu \in \mathcal{P} : \exists \text{ a CM' mechanism } w \text{ on } \Theta^* \text{ s.t.} \\ \mu [\{\theta : U'_i(\theta_i|w) \in (0, 1/n), \forall i \in I\}] > 1 - \frac{1}{n} \end{array} \right\}.$$

Then, we can prove that AFSE' priors are generic in the space of all priors by replacing \mathcal{A}_n with \mathcal{A}'_n and Proposition 2 with Proposition 1 in the proof of Theorem 1; we can also prove that FSE' priors are generic in the space of all finite priors by replacing \mathcal{A}_n^f with $\mathcal{A}'_n \cap \mathcal{P}^f$ in the proof of Theorem 2.

Second, we comment on the issue of IR for AFSE' priors. For an AFSE' prior μ , even if $U'_i(\theta_i|w) \in [0, \varepsilon]$ with probability $1 - \varepsilon$ under w , $U'_i(\theta_i|w)$ may still be negative with positive probability. We can use the idea in Lemma 8 to address this issue. However, it requires some additional restrictions on the mechanism design environment and some careful modification of the definition of AFSE:

- The environment $(X, (u_i)_{i \in I})$ satisfies *excludability* if (i) $X = \prod_{i \in I} X_i$; (ii) u_i does not depend on X_{-i} , i.e., $u_i : X_i \times V \times \mathbb{R} \rightarrow \mathbb{R}$; (iii) for each $i \in I$, there is some $x_i^0 \in X_i$ such that $u_i(x_i^0, v, 0) = 0$ for any $v \in V$.
- μ is an AFSE prior if for any $\varepsilon > 0$, there is an IR and BIC mechanism (q, m) on Θ^μ such that (1) $\mu[\{\theta : U_i(\theta_i|q, m) \in [0, \varepsilon], \forall i\}] \geq 1 - \varepsilon$; (2) $\mu[\{\theta : q_i(\theta) = q'_i(\theta), \forall i\}] \geq 1 - \varepsilon$.

Given excludability, we can follow the proof of Lemma 8 and replace $q_i(\theta) = 0$ with the allocation x_i^0 . As a result, for any AFSE' prior μ and any $\varepsilon > 0$, we can find an IR and BIC mechanism (q, m) such that (1) and (2) in the above definition of AFSE priors hold. That is, given excludability, every AFSE' prior is an AFSE prior. Therefore, the genericity of AFSE priors follows from the genericity of AFSE' priors.

Finally, excludability clearly holds in any private goods allocation environment (e.g., auction, trade, or bargaining) and other environments like regulation or income taxation, but it rules out some prominent environment like public goods provision.

4.2 Implementability

We say a mechanism is value-measurable if the allocations and the payments depend only on the reported values. (Duggan, 1997, Proposition 4 and Theorem 2) proves that every value-measurable mechanism that achieves BIC is virtually Bayesian implementable. In Chen and Xiong (2012), we generalize this result to first-order mechanisms.¹⁸ That is, for any $\varepsilon > 0$ and any mechanism (q, m) that we employ to achieve FSE/AFSE for a prior μ , there exists a mechanism which maps each θ to a random outcome such that in *any* Bayesian Nash equilibrium of the mechanism, the outcome $(q(\theta), m(\theta))$ is obtained with probability $1 - \varepsilon$ for μ -almost all θ .

4.3 Finer Topologies

We prove our topological genericity results under the weak* topology, which is often regarded as a coarse topology. A natural question is whether the results still hold if we endow the space of priors with a finer topology. Recall that our genericity notion is defined using a residual set which is a countable intersection of open and dense sets. A finer topology has more open sets. Hence, the openness in our genericity results (i.e., Propositions 1 and 2) continues to hold in any finer topologies.

For denseness, consider Theorem 1 for example. Recall that Proposition 2 holds for any $\mu \in \mathcal{F}^{cm}$. Hence, as long as \mathcal{F}^{cm} is dense, \mathcal{A}_π defined in (14) still contains an open and dense set. As a result, Theorem 1 still holds with the same proof. Similar ideas work for Theorems 2-4. Using this idea, we prove, in Chen and Xiong (2012), that FSE is generic under the weak* topology combined with the convergence of supports in Hausdorff topology.¹⁹

¹⁸Brusco (1998) constructs an example in which FSE is possible and yet no mechanism can yield FSE as the unique Bayesian Nash equilibrium outcome. That is, FSE need not be exact implementable even when it is possible. In contrast, our result shows that whenever FSE or AFSE is achieved in a BIC first-order mechanism (which we show is possible for generic priors), it is virtually Bayesian implementable.

¹⁹Convergence of priors in the weak* topology need not imply convergence of their supports in the Hausdorff topology. See (Aliprantis and Border, 2006, pp.562-563).

However, we show in [Chen and Xiong \(2012\)](#) that \mathcal{A} is non-generic in \mathcal{P} under the topology induced by the total variation norm. In particular, we further prove that the total variation norm topology is equivalent to the *discrete topology* in the space of finite models. Hence, AFSE/FSE is neither generic nor non-generic in the space of finite models under the total variation norm topology. This suggests that the total variation norm topology is too fine for our purpose. Specifically, under the the total variation norm topology, even CM's genericity result no longer holds.²⁰

4.4 Priors on the universal type space

Throughout the paper, we restrict our attention to priors on the universal type space Θ^* . We provide below a sense in which this is without loss of generality.

Let $(\widehat{\Theta}_i, \widehat{v}_i, \widehat{b}_i)_{i \in I}$ be a (private-value) *abstract type space*, where $\widehat{\Theta}_i$ is a compact metric space of player i 's types; $\widehat{v}_i : \widehat{\Theta}_i \rightarrow \mathbb{R}$ is a continuous function that identifies the value of a type $\widehat{\theta}_i$ being $\widehat{v}_i(\widehat{\theta}_i)$; and $\widehat{b}_i : \widehat{\Theta}_i \rightarrow \Delta(\widehat{\Theta}_{-i})$ is a continuous function that identifies the belief of $\widehat{\theta}_i$ being $\widehat{b}_i(\widehat{\theta}_i)$. Each belief subspace in Θ^* naturally induces an abstract type space, and conversely, an abstract type space can be embedded into the universal type space as a belief subspace in a manner that preserves all the values and beliefs. Formally, let $\eta \equiv (\eta_i)_{i \in N}$ be the canonical embedding from any $\widehat{\Theta}_i$ to Θ_i^* . [Mertens and Zamir \(1985\)](#) and [Heifetz and Neeman \(2006\)](#) show that for each $\widehat{\theta}_i \in \widehat{\Theta}_i$, and any Borel subset E_{-i} of Θ_{-i}^* , we have

$$v_i(\eta_i(\widehat{\theta}_i)) = \widehat{v}_i(\widehat{\theta}_i); \quad (16)$$

$$b_i(\eta_i(\widehat{\theta}_i)) [E_{-i}] = \widehat{b}_i(\widehat{\theta}_i) [\eta_{-i}^{-1}(E_{-i})]. \quad (17)$$

Generally, the existence of a mechanism that achieves FSE on $\widehat{\Theta}$ is not equivalent to the existence of a mechanism that achieves FSE on the belief subspace $\eta(\widehat{\Theta}) \subset \Theta^*$.

²⁰CM fix the number of types for each player and the values associated with the types. In this case, the space of priors is a finite-dimensional simplex, and a natural topology on the space of priors is the Euclidean metric, which coincides with the weak* topology. However, this topology is strictly coarser than the the total variation norm topology. In particular, under the total variation norm topology, a prior which is sufficiently close to a prior which surrenders rent must also surrender rent.

However, for practical reasons, a mechanism designer may be bound to use simple mechanisms. Assume that first-order mechanisms are the only feasible mechanisms. For example, most of the auctions are first-order mechanisms. Under a first-order mechanism, the incentive of a type is fully characterized by her second-order belief. Consequently, the equilibrium outcome of an abstract type space $\hat{\Theta}$ is fully preserved on $\eta(\hat{\Theta})$. This intuition is formalized in the following proposition.

Proposition 3 *For any abstract type space $(\hat{\Theta}_i, \hat{v}_i, \hat{b}_i)_{i \in I}$, there is a first-order mechanism which achieves FSE on $\hat{\Theta}$ if and only if there is a first-order mechanism which achieves FSE on $\eta(\hat{\Theta}) \subset \Theta^*$.*

That is, regarding achieving FSE by first-order mechanisms, any abstract type space $\hat{\Theta}$ can be fully represented by its counterpart $\eta(\hat{\Theta})$ in the universal type space. In this sense, it is without loss of generality to focus on priors on the universal type space. The proof of Proposition 3 is straightforward and it is omitted.

5 Conclusion

In this paper we provide an sense in which full surplus extraction is generically possible even if we relax the common-knowledge assumption of a fixed finite number of types in CM. In other words, private information generically confers no rent to its possessor, whether or not we relax this assumption on the information structures generated by common priors.

As explained in the introduction, the genericity of FSE is an important criterion to evaluate the validity of the classical mechanism design model. Thus, we may have to treat the classical model and its associated theories with caution if we fail to identify inessential assumptions of the classical model which explain the genericity of FSE.

The gist of our analysis is that CM mechanisms are robust to small mis-specifications of priors. This advantage makes it even more puzzling that CM mechanisms prescribed by the classical model are rarely seen in reality. Indeed, as (McAfee and Reny, 1992, p.419)

has argued: "This indicates (at least to us) that the prevalence of the English auction in selling items whose value is uncertain is almost certainly not due to the fact that sellers are maximizing expected revenue." Our results call for further scrutiny of this puzzle.

A Appendix

A.1 Proof of Lemma 8

Lemma 8. *Let w be a CM mechanism defined on Θ^* and (q, m) be a mechanism defined on Θ^* as follows.*

$$(q_i(\theta), m_i(\theta)) = \begin{cases} (q_i^*(\theta), m_i^*(\theta) + w_i(v_{-i}(\theta))), & \text{if } U_i(\theta_i|w) > 0; \\ (0, 0), & \text{if } U_i(\theta_i|w) \leq 0. \end{cases}$$

Then, (q, m) is a first-order mechanism which achieves IR and BIC on Θ^* .

Proof. The value of $U_i(\theta_i|w)$ depends only on the first-order belief of θ_i . Hence, (q, m) is a first-order mechanism.

For any θ_i and θ'_i in Θ_i^* , we have

$$U_i(\theta'_i|\theta_i, q, m) = \begin{cases} U_i(\theta'_i|\theta_i, w), & \text{if } U_i(\theta'_i|w) > 0; \\ 0, & \text{if } U_i(\theta'_i|w) \leq 0. \end{cases} \quad (18)$$

As a result, for every $\theta_i \in \Theta_i^*$, we have

$$U_i(\theta_i|q, m) = \max\{U_i(\theta_i|w), 0\} \geq 0. \quad (19)$$

That is, IR holds. We now check BIC. Pick any $\theta_i \in \Theta_i^*$ and we will show that $U_i(\theta_i|q, m) \geq U_i(\theta'_i|\theta_i, q, m)$ for any $\theta'_i \in \Theta_i^*$. Since w is DSIC and hence BIC on Θ^* , for any possible deviation $\theta'_i \in \Theta_i^*$, we have

$$U_i(\theta_i|w) \geq U_i(\theta'_i|\theta_i, w). \quad (20)$$

There are two cases to check.

Case 1: $U_i(\theta'_i|w) > 0$. Then,

$$U_i(\theta_i|q, m) \geq U_i(\theta_i|w) \geq U_i(\theta'_i|\theta_i, w) = U_i(\theta'_i|\theta_i, q, m),$$

where the first inequality follows from (19); the second inequality follows from (20); the equality follows from (18) and $U_i(\theta'_i|w) > 0$.

Case 2: $U_i(\theta'_i|w) \leq 0$. Then,

$$U_i(\theta_i|q, m) \geq 0 = U_i(\theta'_i|\theta_i, q, m),$$

where the inequality follows from (19) and the equality follows from (18) and $U_i(\theta'_i|w) \leq 0$. Therefore, it is not profitable for type θ_i to deviate to report θ'_i under (q, m) . Hence, BIC is satisfied. ■

A.2 Proof of Lemma 2

Lemma 2. *If $Y (\subset X)$ is dense in X and $U (\subset X)$ is generic in X , then $U \cap Y$ is generic in Y under the relative topology.*

Proof. By the genericity of U , we have $U \supset \bigcap_{n=1}^{\infty} E_n$, where $E_n (\subset X)$ is open and dense in X for every n . Hence,

$$U \cap Y \supset \bigcap_{n=1}^{\infty} (E_n \cap Y). \quad (21)$$

First, $E_n \cap Y$ is open in Y under the relative topology, because E_n is open in X . Second, $E_n \cap Y$ is dense in Y under the relative topology. To see this, for any $y \in Y$ and any open set $F \subset Y$ such that $y \in F$, it suffices to show that there exists some $e \in [E_n \cap Y] \cap F$. Since F is open in Y under the relative topology, we have $F = G \cap Y$, where $G (\subset X)$ is open in X . Since E_n is dense in X , we have $E_n \cap G \neq \emptyset$ and $E_n \cap G$ is open in X . Since Y is dense in X , we have $[E_n \cap G] \cap Y \neq \emptyset$. Thus, there exists $e \in [E_n \cap G] \cap Y$. Finally,

$$[E_n \cap G] \cap Y = [E_n \cap Y] \cap [G \cap Y] = [E_n \cap Y] \cap F.$$

Therefore, there exists some $e \in [E_n \cap Y] \cap F$, and $E_n \cap Y$ is dense in Y under the relative topology. ■

References

- ABREU, D., AND H. MATSUSHIMA (1992): "Virtual Implementation in Iteratively Undominated Strategies: Complete Information," *Econometrica*, 60, 993–1008.
- ALIPRANTIS, C., AND K. BORDER (2006): *Infinite Dimensional Analysis*. Berlin: Springer-Verlag.
- ANDERSON, R., AND W. ZAME (2001): "Genericity with Infinite Many Parameters," *Advances in Theoretical Economics*, 1, 1–62.
- BARELLI, P. (2009): "On the Genericity of Full Surplus Extraction in Mechanism Design," *Journal of Economic Theory*, 144, 1320–1332.
- BERGEMANN, D., AND S. MORRIS (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.
- BRUSCO, S. (1998): "Unique Implementation of the Full Surplus Extraction Outcome in Auctions with Correlated Types," *Journal of Economic Theory*, 80, 185–200.
- CHEN, Y.-C., AND S. XIONG (2011): "The Genericity of Beliefs-Determine-Preferences Models Revisited," *Journal of Economic Theory*, 146, 751–761.
- (2012): "Supplement to 'Genericity and Robustness of Full Surplus Extraction'," mimeo.
- CHRISTENSEN, J. P. R. (1974): *Topology and Borel Structure*, vol. 10 of *North-Holland Mathematical Studies*. North-Holland, Amsterdam.
- CHUNG, K.-S., AND J. C. ELY (2007): "Foundations of Dominant-Strategy Mechanisms," *Review of Economic Studies*, 74, 447–476.
- CRÉMER, J., AND R. MCLEAN (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 53, 345–361.
- DEKEL, E., D. FUDENBERG, AND S. MORRIS (2006): "Topologies on Types," *Theoretical Economics*, 1, 275–309.
- DUDLEY, R. (2002): *Real Analysis and Probability*. Cambridge University Press.

- DUGGAN, J. (1997): "Virtual Bayesian Implementation," *Econometrica*, 65, 1175–1199.
- ELY, J. C., AND M. PEŠKI (2011): "Critical Types," *Review of Economic Studies*, forthcoming.
- HEIFETZ, A., AND Z. NEEMAN (2006): "On the Generic (Im)Possibility of Full Surplus Extraction in Mechanism Design," *Econometrica*, 74, 213–234.
- HUNT, B., T. SAUER, AND J. YORKE (1992): "Prevalence: A Translation-Invariant 'Almost Every' on Infinite-Dimensional Spaces," *Bulletin (New Series) of the American Mathematical Society*, 27, 217–238.
- LAFFONT, J.-J., AND D. MARTIMORT (2000): "Mechanism Design with Collusion and Correlation," *Econometrica*, 68, 309–342.
- MAS-COLELL, A. (1985): *The Theory of General Economic Equilibrium: A Differentiable Approach*, Econometric Society Monographs. Cambridge University Press.
- MCAFEE, R., AND P. RENY (1992): "Correlated Information and Mechanism Design," *Econometrica*, 60, 395–421.
- MERTENS, J.-F., S. SORIN, AND S. ZAMIR (1994): *Repeated games, Part A: Background material*. Discussion Paper 9420, Center for Operations Research and Econometrics, Catholic University of Louvain.
- MERTENS, J.-F., AND S. ZAMIR (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1–29.
- MYERSON, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.
- NEEMAN, Z. (2004): "The Relevance of Private Information in Mechanism Design," *Journal of Economic Theory*, 117, 155–177.
- PETERS, M. (2001): "Surplus Extraction and Competition," *Review of Economic Studies*, 68, 613–631.
- ROBERT, J. (1991): "Continuity in Auction Design," *Journal of Economic Theory*, 55, 169–179.

STINCHCOMBE, M. (2000): "The Gap Between Probability and Prevalence: Loneliness in Vector Spaces," *Proceedings of the American Mathematical Society*, 129, 451–457.