ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 81, No. 2 (March, 2013), 825–847

GENERICITY AND ROBUSTNESS OF FULL SURPLUS EXTRACTION

YI-CHUN CHEN
National University of Singapore, 117570, Singapore

SIYANG XIONG
Rice University, Houston, TX 77005, U.S.A.

Econometrica, Vol. 81, No. 2 (March, 2013), 825–847

GENERICITY AND ROBUSTNESS OF FULL SURPLUS EXTRACTION

BY YI-CHUN CHEN AND SIYANG XIONG[1]

We study whether priors that admit full surplus extraction (FSE) are generic, an issue that 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 made by Crémer 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 *Crémer–McLean mechanisms*.

KEYWORDS: Surplus extraction, information rents, universal type space, common prior, genericity, residual set.

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) showed 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. CM's result thus 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 institutional design" (McAfee and Reny (1992, p. 400)).[2]

[1]We thank the editor and five anonymous referees for insightful comments which have greatly improved the paper. For their helpful comments and discussion, we also thank Tilman Börgers, Eddie Dekel, Jeff Ely, Eduardo Faingold, Songying Fang, Alia Gizatulina, Aviad Heifetz, Martin Hellwig, Jinwoo Kim, Maciej Kotowski, Takashi Kunimoto, Jihong Lee, Qingmin Liu, Xiao Luo, Claudio Mezzetti, Hervé Moulin, Daisuke Oyama, Wolfgang Pesendorfer, Marciano Siniscalchi, Yeneng Sun, Olivier Tercieux, Chih-Chun Yang, and the participants in seminars/conferences at Academia Sinica, Games 2012, the Kansas Workshop on Economic Theory, Rice University, Shanghai University of Finance and Economics, University of Tokyo, and the 11th and the 12th SAET Conferences. We also gratefully acknowledge financial support from National Science Foundation Grant SES-1227620 and Singapore Ministry of Education Academic Research Fund Tier 1. All remaining errors are our own.

[2]CM proved their result in the single-object private-value auction setting with a fixed finite type space. McAfee and Reny (1992) extended CM's result and characterized FSE in a general mechanism design setting with a continuum of types whose beliefs are given by continuous density functions.

 DOI: 10.3982/ECTA10123

Many explanations have been offered to address this *FSE critique*. We now know 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), Che and Kim (2006), and Peters (2001)). Nevertheless, the classical model that incorporates *all* these assumptions is still commonly used. By modifying these essential assumptions of the classical model, the studies just cited demonstrate only 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 provide such an explanation, we must be able to attribute the genericity of FSE to assumptions that are *inessential* to the classical model and yet are critical to the genericity of FSE. Failing to find such inessential assumptions would invite reexamination of the classical model and all the theories based on it.

In an important paper, Heifetz and Neeman (2006; hereafter, HN) identified one such assumption. In particular, HN pointed 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 will need to study the genericity/nongenericity of FSE in the space of general priors supported on an arbitrary number of types. Following this view, HN proved that FSE priors are "negligible" (nongeneric) in both 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)).[3]

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. That is, 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.[4] We say a set is (topologically) generic if it

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

[4]Agents choose their best strategies 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 that makes these "expected values" continuous in beliefs and it is

is a residual set, that is, it contains a countable intersection of open and dense sets; a set is nongeneric if its complement is generic, that is, it is a meager set.[5]

We report two genericity results. We first consider the space of all priors, which can be either finite or infinite. To address the technical difficulties associated with infinite priors, we follow McAfee and Reny (1992) and define a prior to be (almost) FSE if, for any $\varepsilon > 0$, we can construct a mechanism that surrenders at most a surplus of quantity ε to the agents, and moreover, it is Bayesian incentive compatible and individually rational. Our first result is that FSE is generic among all priors.[6]

We then study an important subclass of priors called *models*. A model is a prior that is not a convex combination of other priors. In Section 3.4, we discuss the 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 apply and our genericity result remains the same.

For simplicity, we focus throughout the main text on the single-object, private-value auction environment, and we discuss general mechanism design problems in Section 4.1. Unlike HN, we prove our genericity results by directly studying the mechanism design problems on general priors.[7] We consider a class of mechanisms called *Crémer–McLean mechanisms* (CM mechanisms). A CM mechanism is a second-price auction supplemented with side payments (see Definition 7). The gist of our proof is a robustness property of CM mechanisms: for any $\varepsilon > 0$, if a CM mechanism extracts, for a prior μ, all of the total surplus except for a quantity less than ε, 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).

often regarded as a coarse topology. Barelli (2009) and Chen and Xiong (2011) also adopted 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 took the weak* topology as an important candidate for it. See Section 4.3 for more discussion on other topologies.

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

[6]That is, generically, FSE can be partially implemented. In Chen and Xiong (2013), we further proved that, generically, FSE can be virtually Bayesian fully implemented in the sense of Abreu and Matsushima (1992) and Duggan (1997). We discuss this in Section 4.2.

[7]HN'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 1 to a set of type profiles in which no distinct types have the same belief. We showed in Chen and Xiong (2011) that priors that satisfy the BDP property are topologically generic. Since BDP is necessary but not sufficient for FSE (see HN's Proposition 2), the genericity of BDP priors has no implication for the genericity or nongenericity of FSE priors.

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

Economic modelers obviously would not want to rule out the correlated private information, which admits FSE, 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 imposes a fundamental restriction for economic modelers—revenue maximization aligns with robustness only under correlated information.[8]

Finally, we note that the fact that HN and this paper reach the opposite conclusions is due solely to the different perspectives we take, with HN taking the geometric/measure-theoretical approach while 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, deciding which one is more appropriate depends on the specific context.[9] Despite the difference, our result shows that the measure-theoretic nongenericity of FSE represents a knife-edge situation because, for any $\varepsilon > 0$, the set $\mathcal{F}_\varepsilon$, being open and dense, is not negligible even in the measure-theoretic sense.[10]

[8]The 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 to the classical model, the FSE property is among the persistent (robust) and typical (generic) properties that "have a good chance to be observed."

[9]Anderson and Zame (2001) pointed out some weakness of the residual (meager) set as the notion of genericity (nongenericity). Stinchcombe (2001) discussed some caveats of the prevalence (shy) set as the notion of genericity (nongenericity).

[10]This follows because an open set cannot be finitely shy (see Anderson and Zame (2001)). Thus, this observation holds whenever finite shyness is defined using a topology that is finer than the weak* topology.

The rest of this paper is organized as follows. 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 space of all probability measures on the Borel σ-algebra endowed with the weak* topology. The weak* topology is metrizable under the Prohorov metric, defined as

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

where $E^{\varepsilon} \equiv \{x' : \inf_{x \in E} d_X(x', x) < \varepsilon\}$ (see Dudley (2002, Theorem 11.3.3)). All product spaces are endowed with the product topology, and subspaces are endowed with the relative topology. The support of a probability measure $\mu \in \Delta(X)$, denoted by $\operatorname{Supp}\mu$, is the intersection of all closed sets with measure 1 under μ. For every finite set $F \subset X$, let $|F|$ denote the cardinality of F.

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, a set 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 types of bidder i (see Mertens and Zamir (1985), Heifetz and Neeman (2006, pp. 228–229), and also Section 4.4). Let $v_i : \Theta_i^* \to V_i$ and $b_i : \Theta_i^* \to \Delta(\Theta_{-i}^*)$ 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 every $\theta \in \Theta$, $\{\theta_i\} \times \operatorname{Supp} b_i(\theta)$ is a subset of Θ. For any $\mu \in \Delta(\Theta^*)$, denote 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 every bounded measurable function $\varphi : \Theta^* \to \mathbb{R}$,

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

Let $\mathcal{P} \subset \Delta(\Theta^*)$ be the set of all priors and let $d_{\mathcal{P}}$ denote the Prohorov metric on $\mathcal{P}$. The support of a prior μ is a belief subspace (see Mertens, Sorin, and Zamir (1994, p. 147, item 2)) and is denoted by Θ^μ.[11] A *finite prior* is a prior whose support is a finite set. Let $\mathcal{P}^f$ denote the space of finite priors. A *model* is a prior μ such that there exist no priors π and v and $\alpha \in (0,1)$ such that $\Theta^\pi \neq \Theta^v$ and $\mu = \alpha\pi + (1-\alpha)v$. Let $\mathcal{M}$ denote the space of 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.

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, we can restrict attention to direct mechanisms without loss of generality. A (direct) mechanism on a belief subspace Θ is a list of measurable functions $(q,m) \equiv (q_i : \prod_{j\in I}\Theta_j \to [0,1], m_i : \prod_{j\in I}\Theta_j \to \mathbb{R})_{i\in I}$ satisfying $\sum_{i\in I} q_i(\theta) \leq 1$ for every $\theta \in \prod_{i\in I}\Theta_i$. 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 $\theta_{-i} \in \Theta_{-i}$, that is,

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

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, that is,

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

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

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

[11] In the literature, a belief subspace that 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.

DEFINITION 2: A mechanism (q, m) defined on a belief subspace Θ is (Bayesian) incentive compatible (IC) if

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

That is, (q, m) on Θ satisfies IC iff truthful reporting constitutes a Bayesian Nash equilibrium on Θ.

DEFINITION 3: 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 a surplus of quantity ε is surrendered to the bidders. Following McAfee and Reny (1992, p. 400), we now define (almost) full surplus extraction as follows.

DEFINITION 4: A prior μ is an (almost) *full-surplus-extraction (FSE)* prior if, for any $\varepsilon > 0$, there exists a mechanism (q, m) on Θ^μ that achieves IR, IC, and ε-SE.

We use $\mathcal{F}$ to denote the set of FSE priors in $\mathcal{P}$. 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,

$$\begin{aligned}&b_i^1(\theta_i)[\overline{V}] = b_i^1(\theta_i)\big[\big\{\theta_{-i} \in \Theta_{-i} : \big(v_i(\theta_i), v_{-i}(\theta_{-i})\big) \in \overline{V}\big\}\big],\\ &\quad \forall \text{Borel set } \overline{V} \subset V.\end{aligned}$$

DEFINITION 5: 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

$$\begin{aligned}(4)\qquad &\max_{i\in I} v_i(\theta) = \sum_{i\in I} v_i(\theta_i) q_i^*(\theta_i, \theta_{-i});\\ &m_i^*(\theta) = q_i^*(\theta) \times \max_{j\neq i} v_j(\theta).\end{aligned}$$

It is well known that (q^*, m^*) is IR and IC. Furthermore, (q^*, m^*) is a first-order mechanism. Indeed, (q^*, m^*) depends only on the reported values, that is, 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 $b^1_i(\theta_i) = b^1_i(\theta'_i)$ implies $v_i(\theta_i) = v_i(\theta'_i)$, it follows that (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.

DEFINITION 6: In a topological space X, we say a subset of X is generic if it is a residual set; we say a subset of X is nongeneric if it is a meager set.

The notions of residual sets and meager sets are usually defined and studied in a *Baire space*. 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 thus nonempty) and not meager. Lemma 2 shows that the two spaces of priors $\mathcal{P}$ and $\mathcal{M}$ that we study in this paper are both Baire spaces. Lemma 3 provides a technical result that will be used later. The proofs are relegated to the Appendix.

LEMMA 2: *Both $\mathcal{P}$ and $\mathcal{M}$ are Baire spaces.*

LEMMA 3: *If Y is a dense subset of X and U is generic in X, then $U \cap Y$ is generic in Y.*

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. Second, we establish a robustness property of CM mechanisms (Lemma 8). We then prove that full surplus extraction is generic in the space of all priors (Theorem 1). Finally, we prove that the result also holds for the space of all models (Theorem 2).

3.1. *CM Mechanisms*

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

DEFINITION 7: A mechanism (q, m) defined on Θ^* is called a Crémer–McLean (CM) mechanism if there exists a profile of continuous functions $(w_i : V_{-i} \to \mathbb{R})_{i\in I}$ 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.[12] It has three distinct features. First, w_i depends only on the reported values, which implies that a CM 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 bidder i's opponents 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 with 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. These properties are summarized in the following lemma.

LEMMA 4: *Every CM mechanism is a first-order mechanism and it satisfies IC.*

We now review CM's FSE result as follows. Define

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

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

$$\mathbf{B}^\mu_i \equiv \left[b_i(\theta_i)[\theta_{-i}]\right]_{\theta_i\in\Theta^\mu_i, \theta_{-i}\in\Theta^\mu_{-i}},$$

where each row corresponds to a type θ_i of bidder i and his belief about his opponents' types, that is, $[b_i(\theta_i)[\theta_{-i}]]_{\theta_{-i}\in\Theta^\mu_{-i}}$. We say $\mathbf{B}^\mu_i$ has *full rank* if its column space has rank n, and μ has full rank if $\mathbf{B}^\mu_i$ has full rank for every $i \in I$. Define

$$\mathcal{F}^{\mathrm{cm}} \equiv \big\{\mu \in \mathcal{P} : \exists \text{ a CM mechanism } w \text{ s.t. } U_i(\theta_i|w) = 0,$$
$$\forall (i, \theta) \in I \times \Theta^\mu\big\}.$$

It is straightforward to verify that each μ in $\mathcal{F}^{\mathrm{cm}}$ is an FSE prior.[13] Hereafter, we say μ admits FSE in a CM mechanism if $\mu \in \mathcal{F}^{\mathrm{cm}}$. The FSE result in Crémer and McLean (1988) implies the following lemma.

LEMMA 5: *$\mu \in \mathcal{F}^{\mathrm{cm}}$ if $\mu \in \mathcal{P}^f_n$ has full rank.*

[12]A CM mechanism is defined everywhere on Θ^*. Alternatively, we may define a CM mechanism on a subspace Θ^μ. Or, equivalently, the side payment may be defined only on $v_{-i}(\Theta^\mu_{-i})$ rather than on V_{-i}. Every such CM mechanism corresponds to a CM mechanism in our definition: since $v_{-i}(\Theta^\mu_{-i})$ is a closed set, by the Tietze Extension theorem (see Aliprantis and Border (2006, Theorem 2.47)) for any continuous $w'_i : v_{-i}(\Theta^\mu_{-i}) \to \mathbb{R}$, there exists a continuous function $w_i : V_{-i} \to \mathbb{R}$ such that $w_i(\widehat{v}_{-i}) = w'_i(\widehat{v}_{-i})$ for any $\widehat{v}_{-i} \in v_{-i}(\Theta^\mu_{-i})$.

[13]More precisely, for each $\mu \in \mathcal{F}^{\mathrm{cm}}$, we achieve 0-SE in a CM mechanism.

Clearly, the set of priors that have full rank is open and dense in $\mathcal{P}_n^f$. This is precisely the genericity result in Crémer and McLean (1988). However, as Heifetz and Neeman (2006) argued, 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 assumption about the size of priors' supports. This, in turn, leads us to study the genericity of FSE in $\mathcal{P}$.[14]

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

$$\mathcal{F}^* \equiv \left\{\mu \in \bigcup_{n=1}^{\infty} \mathcal{P}_n^f : \mu \text{ has full rank, full support, and distinct values}\right\}. \tag{5}$$

Thus, $\mathcal{F}^* \subset \mathcal{M}$. Now consider any $\mu \in \mathcal{F}^*$. Since μ has full rank, we can employ the result in Crémer and McLean (1988) to construct a profile of functions $\widetilde{w} = (\widetilde{w}_i : \Theta_{-i}^\mu \to \mathbb{R})_{i\in I}$ that extracts all the surplus. Moreover, since μ has distinct values, $v_{-i}(\theta_{-i}) \neq v_{-i}(\theta'_{-i})$ for any $\theta_{-i}, \theta'_{-i} \in \Theta_{-i}^\mu$ with $\theta_{-i} \neq \theta'_{-i}$. Then, by the Tietze Extension theorem (see Aliprantis and Border (2006, Theorem 2.47)), there exists a CM mechanism $w = (w_i : V_{-i} \to \mathbb{R})_{i\in I}$ such that $w_i(v_{-i}(\theta)) = \widetilde{w}_i(\theta_{-i})$, that is, 0-SE is achieved under w. As a result, $\mathcal{F}^* \subset \mathcal{F}^{\text{cm}}$. We thus have the following result:

LEMMA 6: $\mathcal{F}^* \subset \mathcal{F}^{\text{cm}} \cap \mathcal{M}$.

The next result shows that $\mathcal{F}^*$ is dense in $\mathcal{P}$. Note that Lemmas 6 and 7 imply that $\mathcal{F}^{\text{cm}} \cap \mathcal{M}$ is dense in both $\mathcal{M}$ and $\mathcal{P}$.

LEMMA 7: *$\mathcal{F}^*$ is dense in $\mathcal{P}$.*

Lemma 7 is an immediate consequence of the following facts: (i) $\mathcal{F}^*$ is dense in $\mathcal{P}^f$[15]; (ii) $\mathcal{P}^f$ is dense in $\mathcal{P}$ (Lemma 1).

[14]An alternative way to relax the common-knowledge assumption is to study the genericity of FSE in $\mathcal{P}^f$. However, we can show that $\mathcal{P}^f$ is meager (in $\mathcal{P}^f$), which implies that every subset of $\mathcal{P}^f$ is both meager and residual. That is, $\mathcal{P}^f$ is not a Baire space, and the notion of residual sets is not a sensible notion of genericity in $\mathcal{P}^f$. To see that $\mathcal{P}^f$ is meager, define $\mathcal{P}^n \equiv \{\mu \in \mathcal{P}^f : |\operatorname{supp}\mu| \leq n\}$. Clearly, $\mathcal{P}^f = \bigcup_{n=1}^\infty \mathcal{P}^n$ and $\mathcal{P}^n$ is closed and has no interior in $\mathcal{P}^f$. Hence, $\mathcal{P}^n$ is nowhere dense and $\mathcal{P}^f$ is meager.

[15]To see the idea, consider the case with two bidders. In the following example, $\mu \in \mathcal{P}^f$ is not a model and has neither full rank nor distinct values. However, $\mu_m \in \mathcal{F}^*$ for all large m and $\mu_m \to \mu$

3.2. *Robustness of CM Mechanisms*

For every CM mechanism w defined on Θ^* and every $\varepsilon > 0$, define

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

That is, $\Omega_i(\varepsilon|w)$ is the set of i's types whose interim expected payoffs under the CM mechanism w fall between 0 and ε. The following lemma provides a preliminary robustness property of CM mechanisms, which will be used later to prove our genericity results.

LEMMA 8: *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 every $\varepsilon, \varepsilon' > 0$, $\{\mu \in \mathcal{P} : \mu_i[\Omega_i(\varepsilon|w)] > 1 - \varepsilon', \forall i \in I\}$ is open.*

Lemma 8 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 that retain only a small surplus, it is still the case for any 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 8 is standard and can be found in Chen and Xiong (2013).

Before we prove our genericity result, we present a proposition. The proposition says that if μ admits FSE in a CM mechanism, then we can construct another CM mechanism that extracts most of the surplus in a neighborhood around μ.

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

---

as $m \to \infty$:

μ	$v_2 = 0$	$v_2 = 1/2$	$v_2 = 1$
$v_1 = 0$	$\frac{1}{4}$	0	0
$v_1 = 1$	0	$\frac{3}{8}$	$\frac{3}{8}$

and

μ_m	$v_2 = 0$	$v_2 = 1/2$	$v_2 = 1$
$v_1 = \frac{1}{m}$	$(1 - \frac{1}{m}) \times \frac{1}{4}$	$\frac{1}{16m}$	$\frac{1}{16m}$
$v_1 = 1 - \frac{1}{m}$	$\frac{1}{16m}$	$(1 - \frac{1}{m}) \times \frac{3}{16}$	$(1 + \frac{1}{m}) \times \frac{3}{16}$
$v_1 = 1 - \frac{1}{2m}$	$\frac{1}{16m}$	$(1 + \frac{1}{m}) \times \frac{3}{16}$	$(1 - \frac{1}{m}) \times \frac{3}{16}$

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

$$w_i(v_{-i}) = w'_i(v_{-i}) - \frac{\varepsilon}{2}, \quad \forall (i, v) \in I \times V.$$

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 $\frac{\varepsilon}{2}$. Therefore, w is a CM mechanism such that $U_i(\theta_i|w) = \frac{\varepsilon}{2} \in (0, \varepsilon)$ for every $(i, \theta) \in I \times \Theta^{\mu}$. Hence, $\mu_i[\Omega_i(\varepsilon|w)] = 1$ for every $i \in I$. It therefore follows that $\mu \in \{\mu' \in \mathcal{P} : \mu'_i[\Omega_i(\varepsilon|w)] > 1 - \varepsilon', \forall i \in I\}$. Furthermore, by Lemma 8(b), $\{\mu' \in \mathcal{P} : \mu'_i[\Omega_i(\varepsilon|w)] > 1 - \varepsilon', \forall i \in I\}$ is open. As a result, 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 $i \in I$. *Q.E.D.*

3.3. *Genericity in the Space of Priors*

In this subsection, we prove the genericity of FSE in $\mathcal{P}$. Note that a CM mechanism may not satisfy IR, and we have to deal with this issue in our proof. In Lemma 9, we first provide a way to modify a CM mechanism so as to achieve IR. The proof is relegated to the Appendix.

LEMMA 9: *Let w be a CM mechanism defined on Θ^*, and let (q, m) be a mechanism defined on Θ^*, as follows*:

$$(6) \qquad \big(q_i(\theta), m_i(\theta)\big) = \begin{cases} \big(q_i^*(\theta), m_i^*(\theta) + w_i\big(v_{-i}(\theta)\big)\big), & \text{if } U_i(\theta_i|w) > 0; \\ (0, 0), & \text{if } U_i(\theta_i|w) \le 0. \end{cases}$$

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

Proposition 2 below says that if μ admits FSE in a CM mechanism, we can construct a mechanism that is both IR and IC, and that extracts most of the surplus in a neighborhood around μ.

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

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

$$(7) \qquad \mu'_i\left[\Omega_i\left(\frac{\varepsilon}{4|I|}\middle|w\right)\right] > 1 - \frac{\varepsilon}{4|M||I|}, \quad \forall i \in I,$$

where

$$(8) \qquad M \equiv \max_{\theta_i \in \Theta_i^*}\big|U_i(\theta_i|w)\big| + 1.$$

Applying Lemma 9 to w, we can construct a first-order mechanism (q, m) defined on Θ^* that achieves IR and IC on Θ^*; moreover, by (6), it follows that

$$U_i(\theta_i|q, m) = \max\{U_i(\theta_i|w), 0\}; \tag{9}$$

$$q_i(\theta) = q_i^*(\theta) \quad \text{if} \quad \theta_i \in \Omega_i\left(\frac{\varepsilon}{4|I|}\middle|w\right). \tag{10}$$

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

$$\begin{aligned}
&\int_{\Theta^{\mu'}}\left[\sum_{i\in I} v_i(\theta_i)q_i(\theta_i, \theta_{-i}) - \sum_{i\in I} m_i(\theta)\right]\mu'[d\theta] \\
&= \sum_{i\in I}\int_{\Theta_i^{\mu'}} U_i(\theta_i|q, m)\mu_i'[d\theta_i] \\
&\le \sum_{i\in I}\left[\int_{\Omega_i((\varepsilon/(4|I|))|w)} U_i(\theta_i|w)\mu_i'[d\theta_i]\right. \\
&\quad \left. + \int_{\Theta_i^{\mu'}\setminus\Omega_i((\varepsilon/(4|I|))|w)} \left|U_i(\theta_i|w)\right|\mu_i'[d\theta_i]\right] \\
&\le |I| \times \left[\frac{\varepsilon}{4|I|} + M \times \frac{\varepsilon}{4M|I|}\right] = \frac{\varepsilon}{2},
\end{aligned} \tag{11}$$

where the first inequality follows from (9); the second inequality follows from (7) and (8). Second,

$$\begin{aligned}
&\int_{\Theta^{\mu'}}\left[\max_{i\in I} v_i(\theta) - \sum_{i\in I} v_i(\theta_i)q_i(\theta_i, \theta_{-i})\right]\mu'[d\theta] \\
&= \int_{\Theta^{\mu'}}\left[\sum_{i\in I} v_i(\theta_i)q_i^*(\theta_i, \theta_{-i}) - \sum_{i\in I} v_i(\theta_i)q_i(\theta_i, \theta_{-i})\right]\mu'[d\theta] \\
&= \int_{\Theta^{\mu'}}\left[\sum_{i\in I} v_i(\theta_i)\left[q_i^*(\theta_i, \theta_{-i}) - q_i(\theta_i, \theta_{-i})\right]\right]\mu'[d\theta] \\
&= \sum_{i\in I}\int_{\Theta_i^{\mu'}\setminus\Omega_i((\varepsilon/(4|I|))|w)} v_i(\theta_i)\left[q_i^*(\theta_i, \theta_{-i}) - q_i(\theta_i, \theta_{-i})\right]\mu_i'[d\theta_i] \\
&\le \sum_{i\in I}\mu_i'\left[\Theta_i^{\mu'}\setminus\Omega_i\left(\frac{\varepsilon}{4|I|}\middle|w\right)\right] \\
&\le |I| \times \frac{\varepsilon}{4M|I|} \le \frac{\varepsilon}{4},
\end{aligned} \tag{12}$$

where the first equality follows from (4); the third equality follows from (10); 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 (7); and the last equality follows because $M \geq 1$.

Combining (11) and (12), 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 $\Theta^{\mu'}$. *Q.E.D.*

THEOREM 1: *$\mathcal{F}$ is generic in $\mathcal{P}$. That is, full surplus extraction is generically possible in the space of priors.*

PROOF: Define

$$\mathcal{F}_n \equiv \left\{\mu \in \mathcal{P} : \exists \text{ a first-order mechanism } (q, m) \text{ on } \Theta^{\mu} \text{ that achieves IR, IC, and } \frac{1}{n}\text{-SE}\right\}. \tag{13}$$

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

By Proposition 2, for any $\mu \in \mathcal{F}^{\text{cm}}$, there exist $\gamma_\mu > 0$ and a first-order mechanism (q, m) defined on Θ^* such that, for any μ' with $d_{\mathcal{P}}(\mu', \mu) < \gamma_\mu$, the mechanism (q, m) achieves IR, IC, and $\frac{1}{n}$-SE on Θ^{μ}. That is,

$$\mathcal{B}_{d_{\mathcal{P}}}(\mu, \gamma_\mu) \equiv \{\mu' \in \mathcal{P} : d_{\mathcal{P}}(\mu', \mu) < \gamma_\mu\} \subset \mathcal{F}_n, \quad \forall \mu \in \mathcal{F}^{\text{cm}}.$$

It follows that

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

Since $\mathcal{F}^{\text{cm}}$ is dense in $\mathcal{P}$ by Lemmas 6 and 7, it follows that $\bigcup_{\mu \in \mathcal{F}^{\text{cm}}} \mathcal{B}_{d_{\mathcal{P}}}(\mu, \gamma_\mu)$ is an open and dense set which is contained in $\mathcal{F}_n$. *Q.E.D.*

3.4. *Genericity in the Space of Models*

We now show that our genericity result holds in the space of all models.

THEOREM 2: *$\mathcal{F} \cap \mathcal{M}$ is generic in $\mathcal{M}$. That is, full surplus extraction is generically possible in the space of models.*

Since $\mathcal{M}$ is dense in $\mathcal{P}$ by Lemmas 6 and 7, Theorem 2 follows from Theorem 1 and Lemma 3.

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

$$\mu': \begin{array}{|c|c|c|}\hline & \theta'_2 & \widetilde{\theta}'_2 \\ \hline \theta'_1 & \frac{1}{4} & \frac{1}{4} \\ \hline \widetilde{\theta}'_1 & \frac{1}{4} & \frac{1}{4} \\ \hline \end{array} \qquad \mu'': \begin{array}{|c|c|c|}\hline & \theta''_2 & \widetilde{\theta}''_2 \\ \hline \theta''_1 & \frac{1}{6} & \frac{1}{3} \\ \hline \widetilde{\theta}''_1 & \frac{1}{3} & \frac{1}{6} \\ \hline \end{array}$$

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

Observe that μ^a is not a model because μ' and μ'' have distinct supports.

In general, every prior can be regarded as a convex combination of models.[16] HN interpreted the convex combination as uncertainty faced by the mechanism designer. However, the designer could conceivably resolve this uncertainty by reducing a mechanism design problem associated with a prior to a mechanism design problem associated with models, as follows. First, for each type θ_i, there is a belief subspace Θ which is the minimal belief subspace among those containing θ_i (see Mertens, Sorin, and Zamir (1994, item 2(e), p. 144)), that is, $\theta_i \in \Theta_i$ and there is no belief subspace $\Theta' \subsetneq \Theta$ such that $\theta_i \in \Theta'_i$.

Second, in Chen and Xiong (2013), we proved that a prior μ is a model iff for μ-almost all θ, Θ^μ is the minimal belief subspace containing θ_i for all $i \in I$. Hence, the designer can ask the agents to simultaneously report the minimal belief subspaces containing their actual types. If their reports do not match, the designer levies a large fine. Since for μ-almost all θ, the support of μ is the minimal belief subspace containing θ_i for all i, truth-reporting is an equilibrium.[17] Finally, since every prior is a convex combination of models, the designer can be sure that the reported belief subspace is the support of some model. For

[16] Since each extreme point of the convex space of priors is a model, this claim follows from Choquet's theorem (see Barelli (2009, Section 4)). Using this observation, Barelli (2009) also argued that HN's analysis is not comparable to CM's analysis, since the latter does not rely on this convex combination.

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

such a designer, the genericity issue is "how often" the resulting model permits FSE. Since the notions of genericity that HN adopted require convexity, and since we know of no other notions of genericity that do not require convexity, we can only apply the topological notion of genericity to the space of models.

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 $u_i : X \times V \times \mathbb{R} \to \mathbb{R}$ is a quasilinear ex post utility of agent i. Consider an arbitrary IC 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

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

We consider the mechanism supplemented with the side-payment scheme $(w_i)_{i\in I}$to be a CM′ mechanism, which we also denote 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, that is, $U_i'(\theta_i) = U_i(\theta_i|q^*, m^*)$ and $U_i'(\theta_i|w) = U_i(\theta_i|w)$. Furthermore, every CM′ mechanism is IC.

In this setup, we first define the notion of FSE′ priors as follows:

- μ is an FSE′ prior if, for any $\varepsilon > 0$, there is a CM′ mechanism w on Θ^μ such that $\mu[\{\theta : U_i'(\theta_i|w) \in (0, \varepsilon), \forall i\}] > 1 - \varepsilon$.

Note that IR holds only with probability arbitrarily close to 1 for an FSE′ 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 FSE′ 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, after replacing $U_i(\theta_i|w)$ with $U_i'(\theta_i|w)$ and CM mechanisms with CM′ mechanisms everywhere in Section 3.2, we proceed as follows. The continuity of U_i' and w_i implies that $U_i'(\theta_i|w)$ is continuous and hence the thesis of Lemma 8. Proposition 1 then follows from Lemma 8. Now define

$$\mathcal{F}_n' \equiv \Big\{\mu \in \mathcal{P} : \exists \text{ a CM}' \text{ mechanism } w \text{ on } \Theta^* \text{ s.t.}$$
$$\mu\big[\big\{\theta : U_i'(\theta_i|w) \in (0, 1/n), \forall i \in I\big\}\big] > 1 - \frac{1}{n}\Big\}.$$

We can then prove that FSE′ priors are generic in the space of all priors by replacing $\mathcal{F}_n$ with $\mathcal{F}'_n$, and Proposition 2 with Proposition 1, in the proof of Theorem 1.

Second, we comment on the issue of IR for FSE′ priors. For an FSE′ 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 address this issue by using the idea in Lemma 9. However, we will need some additional restrictions on the mechanism design environment and some careful modification of the definition of FSE, as shown here:

- 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., $u_i : X_i \times V \times \mathbb{R} \to \mathbb{R}$); (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 FSE prior with respect to (q', m') if, for any $\varepsilon > 0$, there is an IR and IC mechanism (q, m) on Θ^μ such that (1) $\mu[\{\theta : U_i(\theta_i|q, m) \in (0, \varepsilon), \forall i\}] \geq 1-\varepsilon$; and (2) $\mu[\{\theta : q_i(\theta) = q'_i(\theta), \forall i\}] \geq 1-\varepsilon$.

In other words, excludability says that it is possible to "exclude" an agent by assigning him an allocation that generates utility zero (i.e., his reservation utility) regardless of the actual value profile v. Given excludability, we can follow the proof of Lemma 9 and replace $q_i(\theta) = 0$ with the allocation x_i^0 whenever IR for θ_i is violated. We can similarly show that the modified mechanism satisfies IC. As a result, for any FSE′ prior μ and any $\varepsilon > 0$, since IR is violated with only μ-probability ε, we can find an IR and IC mechanism (q, m) such that (1) and (2) in the above definition of FSE priors hold. That is, given excludability, every FSE′ prior is an FSE prior. Therefore, the genericity of FSE priors follows from the genericity of FSE′ priors.

Finally, excludability clearly holds in any private goods allocation environment (e.g., auction, trade, or bargaining) and other environments such as regulation or income taxation, but it rules out some prominent environments such as public goods provision.[18]

4.2. *Implementability*

We say that a mechanism is value-measurable if the allocations and the payments depend only on the reported values. By Duggan (1997, Proposition 4 and Theorem 2), every value-measurable mechanism that achieves IC is virtually Bayesian implementable. In Chen and Xiong (2013), we generalized this result to first-order mechanisms.[19] That is, for any $\varepsilon > 0$ and any mechanism

[18]We also note that the IR constraint is not always relevant. For example, Fudenberg and Tirole (1991, p. 245) noted that "in some public good problems, the government may impose decisions that the agents cannot veto" and "whether an individual rationality constraint should be included in the model depends on the extent of the coercive power of the principal, or equivalently, on the distribution of property rights."

[19]Brusco (1998) constructed 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 Bayesian

(q, m) that we employ to achieve FSE for a prior μ, there exists a mechanism that 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. Since a finer topology has more open sets, 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 $\mu \in \mathcal{F}^{\mathrm{cm}}$. Hence, as long as $\mathcal{F}^{\mathrm{cm}}$ is dense, $\mathcal{F}_n$, defined as in (13), still contains an open and dense set. As a result, Theorem 1 still holds with the same proof. Using this idea, we proved in Chen and Xiong (2013) that our genericity results remain true under the weak* topology combined with the convergence of supports in the Hausdorff topology.[20]

However, we also showed in Chen and Xiong (2013) that $\mathcal{F}$ is nongeneric in $\mathcal{P}$ under the topology induced by the total variation norm. Furthermore, we proved that the total-variation-norm topology is equivalent to the *discrete topology* in the space of finite models. Hence, FSE is neither generic nor nongeneric 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. In particular, under the total-variation-norm topology, even CM's genericity result no longer holds.[21]

4.4. *Priors on the Universal Type Space*

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

fully implementable even when it is possible. In contrast, our result shows that whenever FSE is achieved in a BIC first-order mechanism (which we show is possible for generic priors), it is virtually Bayesian fully implementable.

[20]Convergence 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).

[21]CM fixed the number of types for each player and the values associated with the types. In this case, since the space of priors is a finite-dimensional simplex, a natural topology on the space of priors is the Euclidean metric topology, which coincides with the weak* topology. However, this topology is strictly coarser than the total-variation-norm topology. In particular, under the total-variation-norm topology, a prior that is sufficiently close to a non-FSE prior must also be a non-FSE prior.

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 bidder i's types; $\widehat{v}_i : \widehat{\Theta}_i \to V_i$ is a continuous function that identifies the value of a type $\widehat{\theta}_i$ being $\widehat{v}_i(\widehat{\theta}_i)$; and $\widehat{b}_i : \widehat{\Theta}_i \to \Delta(\widehat{\Theta}_{-i})$ 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) showed that, for each $\widehat{\theta}_i \in \widehat{\Theta}_i$, and for any Borel subset E_{-i} of Θ^*_{-i}, we have

$$v_i\big(\eta_i(\widehat{\theta}_i)\big) = \widehat{v}_i(\widehat{\theta}_i); \tag{14}$$

$$b_i\big(\eta_i(\widehat{\theta}_i)\big)[E_{-i}] = \widehat{b}_i(\widehat{\theta}_i)\big[\eta^{-1}_{-i}(E_{-i})\big]. \tag{15}$$

Generally, the existence of a mechanism that achieves FSE on $\widehat{\Theta}$ does not imply the existence of a mechanism that achieves FSE on the belief subspace $\eta(\widehat{\Theta}) \subset \Theta^*$. However, for practical reasons, a mechanism designer may be obliged to use simple mechanisms. Assume that first-order mechanisms are the only feasible mechanisms. For example, most auctions are first-order mechanisms. Under a first-order mechanism, the incentive of a type is fully characterized by his second-order belief. Consequently, the equilibrium outcome of an abstract type space $\widehat{\Theta}$ is fully preserved on $\eta(\widehat{\Theta})$. This intuition is formalized in the following proposition.

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

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

5. CONCLUSION

In this paper, we provide a sense in which full surplus extraction is generically possible even if we relax CM's common-knowledge assumption of a fixed finite number of types. In other words, private information generically confers no rent on 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 for evaluating 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 misspecifications of priors. This advantage makes it even more puzzling that CM mechanisms are rarely seen in reality. Indeed, as McAfee and Reny (1992, p. 419) have 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.

APPENDIX: PROOFS OF LEMMAS 2, 3, AND 9

PROOF OF LEMMA 2: First, $\mathcal{P}$ is a compact metric space (see Mertens, Sorin, and Zamir (1994, p. 147, item 2)) and thus a complete metric space. This implies that $\mathcal{P}$ is a Baire space (see Willard (1970, Corollary 25.4)). To see that $\mathcal{M}$ is a Baire space, note first that the set of extreme points, denoted by $\mathcal{M}'$,[22] in the convex compact metric space $\mathcal{P}$ is a G_δ set (see Phelps (2001, Proposition 1.3)) and thus is also a Baire space (see Willard (1970, Theorem 25.3)). Clearly, $\mathcal{F}^* \subset \mathcal{M}'$, and thus it follows from Lemmas 6 and 7 that $\mathcal{M}'$ is dense in $\mathcal{M}$. Since $\mathcal{M}$ contains $\mathcal{M}'$ as a dense subset, it follows that $\mathcal{M}$ is also a Baire space (see Engelking (1989, p. 201, Exercise 3.9.J.(b))). *Q.E.D.*

PROOF OF LEMMA 3: 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,

$$U \cap Y \supset \bigcap_{n=1}^{\infty} (E_n \cap Y). \tag{16}$$

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 $G_y \subset Y$ such that $y \in G_y$, it suffices to show that there exists some $y_n \in [E_n \cap Y] \cap G_y$. Since G_y is open in Y under the relative topology, we have $G_y = G \cap Y$, where G ($\subset X$) is open in X. Since E_n is open and 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 $y_n \in [E_n \cap G] \cap Y$. Finally,

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

Therefore, there exists some $y_n \in [E_n \cap Y] \cap G_y$, and $E_n \cap Y$ is dense in Y under the relative topology. It then follows from (16) that $U \cap Y$ *is generic in* Y. *Q.E.D.*

[22] $\mathcal{M}' \equiv \{\mu \in \mathcal{P}$: there exist no priors π, v, and $\alpha \in (0, 1)$ such that $\pi \neq v$ and $\mu = \alpha\pi + (1 - \alpha)v\}$.

PROOF OF LEMMA 9: 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

$$U_i\big(\theta_i'|\theta_i, q, m\big) = \begin{cases} U_i\big(\theta_i'|\theta_i, w\big), & \text{if } U_i\big(\theta_i'|w\big) > 0; \\ 0, & \text{if } U_i\big(\theta_i'|w\big) \leq 0. \end{cases} \tag{17}$$

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

$$U_i(\theta_i|q, m) = \max\big\{U_i(\theta_i|w), 0\big\} \geq 0. \tag{18}$$

That is, IR holds. We now check IC. For any $\theta_i \in \Theta_i^*$, we 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 satisfies IC on Θ^*, then for any possible deviation $\theta_i' \in \Theta_i^*$, we have

$$U_i(\theta_i|w) \geq U_i\big(\theta_i'|\theta_i, w\big). \tag{19}$$

There are two cases to check.

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

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

where the first inequality follows from (18); the second inequality follows from (19); and the equality follows from (17) and $U_i(\theta_i'|w) > 0$.

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

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

where the inequality follows from (18) and the equality follows from (17) and $U_i(\theta_i'|w) \leq 0$. Therefore, it is not profitable for type θ_i to deviate by reporting θ_i' under (q, m). Hence, IC is satisfied. *Q.E.D.*

REFERENCES

ABREU, D., AND H. MATSUSHIMA (1992): "Virtual Implementation in Iteratively Undominated Strategies: Complete Information," *Econometrica*, 60, 993–1008. [827]

ALIPRANTIS, C., AND K. BORDER (2006): *Infinite Dimensional Analysis*. Berlin: Springer. [833, 834,842]

ANDERSON, R., AND W. ZAME (2001): "Genericity With Infinite Many Parameters," *Advances in Theoretical Economics*, 1, 1–62. [826,828]

BARELLI, P. (2009): "On the Genericity of Full Surplus Extraction in Mechanism Design," *Journal of Economic Theory*, 144, 1320–1332. [827,839]

BRUSCO, S. (1998): "Unique Implementation of the Full Surplus Extraction Outcome in Auctions With Correlated Types," *Journal of Economic Theory*, 80, 185–200. [841]

CHE, Y.-K., AND J. KIM (2006): "Robustly Collusion-Proof Mechanisms," *Econometrica*, 74, 1063–1107. [826]

CHEN, Y.-C., AND S. XIONG (2011): "The Genericity of Beliefs-Determine-Preferences Models Revisited," *Journal of Economic Theory*, 146, 751–761. [827]

——— (2013): "Supplement to 'Genericity and Robustness of Full Surplus Extraction'," *Econometrica Supplementary Material*, 81, http://www.econometricsociety.org/ecta/Supmat/10123_proofs.pdf. [827,835,839,841,842]

CHRISTENSEN, J. P. R. (1974): *Topology and Borel Structure*. North-Holland Mathematical Studies, Vol. 10. Amsterdam: North-Holland. [826]

CHUNG, K.-S., AND J. C. ELY (2007): "Foundations of Dominant-Strategy Mechanisms," *Review of Economic Studies*, 74, 447–476. [839]

CRÉMER, J., AND R. MCLEAN (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257. [825,833,834]

DEKEL, E., D. FUDENBERG, AND S. MORRIS (2006): "Topologies on Types," *Theoretical Economics*, 1, 275–309. [827]

DUDLEY, R. (2002): *Real Analysis and Probability*. Cambridge: Cambridge University Press. [829]

DUGGAN, J. (1997): "Virtual Bayesian Implementation," *Econometrica*, 65, 1175–1199. [827,841]

ELY, J. C., AND M. PĘSKI (2011): "Critical Types," *Review of Economic Studies*, 78, 907–937. [827]

ENGELKING, R. (1989): *General Topology*. Berlin: Heldermann. [844]

FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. Cambridge: MIT Press. [841]

HEIFETZ, A., AND Z. NEEMAN (2006): "On the Generic (Im)Possibility of Full Surplus Extraction in Mechanism Design," *Econometrica*, 74, 213–233. [825,826,829,834,843]

HUNT, B., T. SAUER, AND J. YORKE (1992): "Prevalence: A Translation-Invariant 'Almost Every' on Infinite-Dimensional Spaces," *Bulletin of the American Mathematical Society (New Series)*, 27, 217–238. [826]

LAFFONT, J.-J., AND D. MARTIMORT (2000): "Mechanism Design With Collusion and Correlation," *Econometrica*, 68, 309–342. [826]

MAS-COLELL, A. (1985): *The Theory of General Economic Equilibrium: A Differentiable Approach*. Econometric Society Monographs. Cambridge: Cambridge University Press. [828]

MCAFEE, R., AND P. RENY (1992): "Correlated Information and Mechanism Design," *Econometrica*, 60, 395–421. [825,827,831,840,844]

MERTENS, J.-F., AND S. ZAMIR (1985): "Formulation of Bayesian Analysis for Games With Incomplete Information," *International Journal of Game Theory*, 14, 1–29. [826,829,843]

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. [826,830,839,844]

MYERSON, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73. [825]

NEEMAN, Z. (2004): "The Relevance of Private Information in Mechanism Design," *Journal of Economic Theory*, 117, 55–77. [827]

PETERS, M. (2001): "Surplus Extraction and Competition," *Review of Economic Studies*, 68, 613–631. [826]

PHELPS, R. (2001): *Lectures on Choquet's Theorem*. New York: Springer. [844]

ROBERT, J. (1991): "Continuity in Auction Design," *Journal of Economic Theory*, 55, 169–179. [826]

STINCHCOMBE, M. (2001): "The Gap Between Probability and Prevalence: Loneliness in Vector Spaces," *Proceedings of the American Mathematical Society*, 129, 451–457. [828]

WILLARD, S. (1970): *General Topology*. Reading: Addison-Wesley. [844]

Dept. of Economics, National University of Singapore, 117570, Singapore; ecsycc@nus.edu.sg

and

Dept. of Economics, Rice University, Houston, TX 77005, U.S.A.; xiong@rice.edu.

Manuscript received June, 2011; final revision received October, 2012.