

Nehring, Klaus

Working Paper

Incentive-Compatibility in Large Games

Working Paper, No. 95-16

Provided in Cooperation with:

University of California Davis, Department of Economics

Suggested Citation: Nehring, Klaus (1995) : Incentive-Compatibility in Large Games, Working Paper, No. 95-16, University of California, Department of Economics, Davis, CA

This Version is available at:

<https://hdl.handle.net/10419/189441>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

**INCENTIVE-COMPATIBILITY
IN LARGE GAMES**

Klaus Nehring

Working Paper Series #95-16

Department of Economics
University of California
Davis, California 95616-8578

Incentive-Compatibility in Large Games

Klaus Nehring

Working Paper Series No. 95-16
May 1995

Note: *The Working Papers of the Department of Economics, University of California, Davis, are preliminary materials circulated to invite discussion and critical comment. These papers may be freely circulated but to protect their tentative character they are not to be quoted without the permission of the author.*

SUMMARY

We argue that large games are of analytical interest partly because they can be understood in terms of a unifying condition of incentive-compatibility, strategy-proofness. In contrast to finite games, strategy-proofness applies not only to dominant-strategy equilibria, but also to a large class of Nash equilibria and to Bayesian **Nash** equilibria with independent types.

Based on Kolmogorov's zero-one law, it is also shown that Bayesian Nash equilibria *coincide* with a class of Nash equilibria in games of incomplete information when there is a countably infinite number of players and types are independent.

1. INTRODUCTION

A non-cooperative game is "large" if any single player's actions have only negligible influence on any other player's outcome. This paper argues that large games are of analytical interest partly because they can be understood in terms of a unifying condition of incentive-compatibility, strategy-proofness. In contrast to finite games, strategy-proofness applies not only to dominant-strategy equilibria, but also to a large class of Nash equilibria and to Bayesian Nash equilibria with independent types.

For Nash equilibria in symmetric mechanisms, this has been realized before in the context of private goods continuum economies in the observation that equilibrium outcomes must yield "envy-free" allocations (Hammond (1979), Champsaur-Laroque (1981)). Section 2 of this paper generalizes this insight to non-symmetric mechanisms and strategy-functions satisfying an aggregation condition called "local separability".

Sections 3 and 4 generalize the "revelation theorem for non-atomic games" of section 2 to mixed strategies and to dynamic contexts. The former relies on a law of large numbers argument; the latter extension is worth to be spelled out explicitly as it involves an additional assumption on behavioral strategies.

The mathematical meat of the paper is in section 5 which considers situations of incomplete information. It shows under regularity conditions that if types are independently distributed, Bayesian Nash equilibria are Nash almost-surely; a converse is shown to hold as well. By consequence, Bayesian incentive-compatibility reduces to strategy-proofness. The proof employs one of the more intriguing results of probability theory, Kolmogorov's zero-one law.

Appeals to laws of large numbers in the context of large games under uncertainty are not new; see in particular the recent contributions by Al-Najjar (1994) and Mas-Colell-Vives (1993). New, however, appear to be their application to non-symmetric

mechanisms as well as the **equivalence** result itself. To establish this result in the general case, one needs (and must make do with) a "non-constructive" law asserting the effective elimination of relevant uncertainty by way of aggregation, without being given an explicit specification of the resulting certainty. This is exactly what Kolmogorov's zero-one law delivers.

To make its application possible, section 5 assumes the set of players to be countably infinite. This keeps the mathematics comparatively basic and transparent. Such transparency seems to be of value given that the boundary between a merely technical and a substantive assumption is somewhat blurred in this context; in particular, the key mathematical fact underlying the equivalence result, lemma 2, hinges critically on the measurability of some function. For further discussion of the assumption of countability, the reader is referred to the concluding paragraphs of section 5.

2. A REVELATION THEOREM FOR LARGE GAMES

Let \mathbf{I} be a set of players; in general, \mathbf{I} is to be understood as infinite.

A player i is described by a **type** (= set of characteristics) $\theta_i \in \Theta_i$, with associated strict preference relation \succ_{θ_i} ; i plays a **strategy** $s_i \in \Sigma_i$ and receives an **outcome** $x_i \in X_i$.

Profiles of types ("environments"), strategies, outcomes and the product spaces they belong to are denoted by dropping subscripts.

Let $\theta \sim \theta'$ iff θ and θ' differ in at most a finite numbers of individuals, i.e. if $\#\{i \in I \mid \theta_i \neq \theta'_i\} < \infty$. $\mathcal{D} \subset \Theta$ is *locally closed* if $\theta \in \mathcal{D}$ and $\theta' \sim \theta$ imply $\theta' \in \mathcal{D}$. A locally closed \mathcal{D} is a *local domain* if $\theta, \theta' \in \mathcal{D}$ imply $\theta \sim \theta'$. Local domains are the equivalence classes $[\theta]$ defined by \sim . If \mathbf{I} is infinite, environments in the same local domain are only microscopically different, and \sim can be interpreted as relation of "macroscopic equivalence".

A **mechanism** $h : \Sigma \rightarrow X$ maps strategy profiles to outcome-profiles. A **strategy function profile** $a : \mathcal{D} \rightarrow \Sigma$ maps environments to strategy profiles. Lastly, a **choice function** $C : \mathcal{D} \rightarrow X$ maps environments to outcome profiles.

The mechanism h is **locally separable** if $s_i = s'_i$ and $s \sim s'$ imply $h_i(s) = h_i(s')$. Local separability of strategy- and choice functions (and similar entities) is defined analogously.

Local separability of a mechanism or a choice-function can be viewed as an aggregation property. It seems rather weak as an assumption on a choice function, but not as an assumption on a mechanism. In particular, it precludes all mechanisms that rely on the collective policing of individual agents' truthful reporting of the environment; mechanisms of this kind play a key role in the implementation literature with a finite number of agents, following the lead of Maskin (1985). Considering games with an infinite number of players allows thus to draw a distinction between two types of implementation – locally separable Nash or simple Nash – whose finite equivalents deserve to be worked out by future research.

Local separability of a strategy-function amounts to the effective irrelevance of "microscopic information" in determining an agent's strategic behavior; such irrelevance is the natural result of agents' Nash behavior in a separable mechanism (see remark 1 below). The non-reliance on microscopic information may be an attractive decentralization property of a locally separable mechanism if agents can be assumed to know at most "macroscopic facts" about the environment such as the distribution of other agents' types. From a mechanism-design perspective, this consideration may motivate the design of a locally separable mechanism which makes such information irrelevant.

The strategic behavior described by a induces the choice-function $C = h \circ a$ in the mechanism h . Conversely, any function C defined on the universe of environments \mathcal{O} can be viewed as a "direct" or "revelation" mechanism with $\Sigma_i = \mathcal{O}$; the identity-

function $h \circ \sigma$ describes truthful revelation of types in this case. It is clear that

Observation 1 $h \circ \sigma$ is locally separable whenever both the mechanism and each strategy-function are locally separable.

C is **strategy-proof** if for no $i \in \mathbf{I}$ and no $\theta \in \mathcal{D}$, $\theta'_i \in \Theta_i$:

$C((\theta'_i, \theta_{-i}))(i) \succ_{\theta} C(\theta)(i)$. The strategy-function profile a is in **Nash equilibrium** if, for no $i \in \mathbf{I}$, $\theta \in \mathcal{D}$ and $s'_i \in \Sigma_i$: $h(s'_i, \sigma_{-i}(\theta))(i) \succ_{\theta} h(\sigma(\theta))$.

Remark 1: Note that the best-reply and Nash correspondences are locally separable, whenever h is. Local separability of strategy-functions is thus a natural result of their being in Nash equilibrium. Conversely, it is clear that to ensure the local separability of strategy-functions in Nash-equilibrium, the mechanism itself needs to be locally separable in general.

Note also that, for locally separable mechanisms, the existence of locally separable strategy functions in Nash equilibrium follows easily from the existence of Nash equilibria. Moreover, if best replies are unique, any Nash equilibrium s in the game (h, θ) is uniquely extended to a locally separable strategy function in Nash equilibrium.

One easily verifies:

Proposition 1 Let $h : \Sigma \rightarrow X$ be a locally separable mechanism and $a : \mathcal{D} \rightarrow \Sigma$ locally separable profile of strategy-functions in Nash equilibrium. Then $h \circ a$ is locally separable and strategy-proof.

Proof. From a being in Nash equilibrium, it follows directly that

$$\text{for no } i \in \mathbf{I}, \theta \in \mathcal{D} \text{ and } \theta'_i \in \Theta_i : h(\sigma_i(\theta'_i), \sigma_{-i}(\theta))(i) \succ_{\theta} h(\sigma(\theta)).$$

By the local separability of σ , this implies that

$$\text{for no } i \in \mathbf{I}, \theta \in \mathcal{D} \text{ and } \theta'_i \in \Theta_i : h(\sigma(\theta'_i, \theta_{-i}))(i) \succ_{\theta} h(\sigma(\theta)),$$

i.e. that $h \circ a$ is strategy-proof. (Note that the local separability of h has not been used so far.)

The local separability of $h \circ a$ has been stated in observation 1. \square

The following observation shows that for finite I , the proposition reduces to the standard revelation-theorem for dominant-strategy equilibria (see Gibbard (1973)) after the elimination of unplayed strategies.

Observation 2 Let $h : \Sigma \rightarrow X$ be a locally separable mechanism and $a : \mathcal{D} \rightarrow \Sigma$ locally separable profile of strategy-functions in Nash equilibrium. Let h' denote the restriction of h to $\prod_{i \in I} \sigma_i(\Theta_i)$. If \mathcal{D} is a local domain, in particular, if I is finite, the strategy-profile $\sigma(\theta)$ is a dominant-strategy equilibrium in the game (h', θ) , for any $\theta \in \mathcal{D}$.

Proposition 1 is interesting because for mechanisms operating on non-local domains, dominant-strategy implementation can be achieved only through a genuine redesigning of the strategy-space. In typical "strategic market games" (see, e.g., Dubey-Mas Colell-Shubik (1980)), the strategy space and mechanism are quite "simple", while the equilibrium strategy-functions are "complex" in the sense that agents' strategy choices depend on the knowledge of other agents' characteristics. This "invisible hand" operating through players' beliefs can be replaced by the "visible hand" of the direct mechanism $h \circ a$ which typically is more complex by encoding player's equilibrium computations in the mechanism itself. On the other hand, the strategy functions are simplified, consisting in the revelation of players' own preferences .

Another way to appreciate the strength of the result is to consider it from a mechanism-design perspective. Proposition 1 implies, then, that a planner is constrained to implement strategy-proof choice-functions, even if he possesses aggregate knowledge of the economy such as the joint distribution of agents' preferences

and publicly observable characteristics and if he has the power to recommend (self-enforcing) strategies, as long as he fails to know who has which preference (as long as there is "assignment uncertainty", in the words of Roberts (1984)).

In contrast to Schmeidler's (1973) distributional approach, "effectively large" games have been defined here by conditions of "local separability" without reference to a measure space of agents. Of course, measure-theoretic structure could be added to the model; however, while little is gained for the present analysis, such structure is inconveniently restrictive in the incomplete-information context of section 5.

At a conceptual level, the present approach assumes that individual players and their strategy choices are meaningfully defined entities, even when they are individually negligible in the aggregate. This allows to distinguish statements holding for *all* players from those holding for merely *almost all* players. Nehring (1995) shows in the context of incentive-compatible resource allocation in private-goods economies that such a distinction may make a difference even in the aggregate, and that, moreover, that difference has a counterpart in large finite economies.

3. EXTENSION TO MIXED STRATEGIES

It is now shown how proposition 1 can be applied to Nash equilibria in mixed-strategies by interpreting strategies (or outcomes) as probability-mixtures of pure strategies (respectively outcomes:).

Assume \mathbf{I} to be countably infinite, and, for simplicity, all Σ_i , X_i , to be finite with associated σ -algebras equal to the power sets $\mathcal{P}(\Sigma_i)$ and $\mathcal{P}(X_i)$, and assume h to be measurable (in the natural way); a more general treatment along the lines of section 5 dealing with Bayesian Nash equilibria is of course possible.

The key is the observation that with locally separable mechanisms, an infinite numbers of players effectively eliminates the strategic uncertainty about others' play;

as a result, independently mixed strategies lead to *independent* mixtures of outcomes.

Let $\mathbf{M}(\Sigma_i)$ and $\mathbf{M}(X_i)$ the set of probability measures on $(\Sigma_i, \mathcal{P}(\Sigma_i))$ and $(X_i, \mathcal{P}(X_i))$ respectively. A **mixed strategy profile** $(\mu_i)_{i \in I}$ is an element of $\prod_{i \in I} \mathcal{M}(\Sigma_i)$ and is understood to induce the product measure $\prod_{i \in I} \mu_i$ on the product σ -algebra \mathbf{S} generated by the $\{\mathcal{P}(\Sigma_i)\}_{i \in I}$

Lemma 1 *If h is measurably locally separable and μ is any mixed-strategy profile, there exist functions $\bar{h}_i : \Sigma_i \rightarrow X_i$, for $i \in I$, such that $\bar{h}_i(t) = h(t, s_{-i})$ $(\prod_{i \in I} \mu_i)$ -almost surely, for all $t \in \Sigma_i$.*

Moreover, $(\prod_{i \in I} \mu_i) \circ h^{-1} = \prod_{i \in I} (\mu_i \circ \bar{h}_i^{-1})$.

Proof. The second part of the lemma is a direct consequence of the first.

The first part follows from observing that, for all $i \in I, t \in \Sigma_i$ and $x \in A_i$ the set $T = \{s \mid h(t, s_{-i})(i) = x\}$ is a tail-event in the product-algebra \mathbf{S} .

By Kolmogorov's zero-one law, the probability of T with respect to $\prod_{i \in I} \mu_i$ is zero or one. \square

The lemma implies that the locally separable deterministic mechanism h induces the function $\hat{h} : \prod_{i \in I} \mathbf{M}(\Sigma_i) \rightarrow \prod_{i \in I} \mathbf{M}(X_i)$ with $\hat{h}(\mu) = (\mu_i \circ \bar{h}_i^{-1})_{i \in I}$. Formally, \hat{h} is a mechanism in the sense of section 2 as well, mapping mixed strategies to independent mixtures of outcomes; it evidently inherits the local separability of h . As a result, proposition 1 can be applied. Note that the equivalent direct mechanism (viz. the induced choice-function) of a mixed-strategy equilibrium in a deterministic mechanism is non-deterministic in general by mapping to independent mixtures of outcomes; it is, however, deterministic in the aggregate.

4. EXTENSION TO EXTENSIVE FORM GAMES

It is worthwhile to briefly consider how proposition 1 extends to dynamic settings, since this turns out to involve assumptions on the strategies themselves.

For simplicity, I will consider only extensive games in which temporal strategy-sets do not depend on **history**; moreover, it is not necessary to explicitly specify the information structure of the game. Thus, at each point in time, $t = 1, \dots, T (\leq \infty)$, players make moves $m_i^t \in M_i^t$. A **play** for player **i** a total history of moves $m_i = (m_i^t)_{t \leq T} \in M_i$.

A **behavioral strategy** $s_i = (s_i^t)_{t \leq T}$ maps histories until $t-1$ to moves of **i** at t , i.e. $s_i^t: \prod_{j \in I, j \neq i, \tau < t} M_j^\tau \rightarrow M_i^t$ for all $t \leq T$; let Σ_i denote their set.

Let $\varphi: \prod_{i \in I} \Sigma_i \rightarrow \prod_{i \in I} M_i$ denote the canonical mapping from behavioral strategies to induced plays. In an extensive game context, a mechanism h maps *plays* to outcomes, $h: \prod_{i \in I} M_i \rightarrow \prod_{i \in I} X_i$. h induces a normal-form mechanism $h \circ \varphi: \prod_{i \in I} \Sigma_i \rightarrow \prod_{i \in I} X_i$. To apply proposition 1, one needs to ensure that the induced normal-form mechanism is locally separable; this requires the following assumption on behavioral strategies.

A behavioral strategy s_i is **locally separable** if each s_i^t is, as a function of previous history. Local separability of a behavioral strategy may arise in two ways. First, it may be implied by the information structure if players have only macroscopic information (about \sim -equivalence classes) about others' play. Moreover, a player may simply not bother about the play of a finite number of other players; this would seem quite reasonable in the context of a locally separable mechanism h , since he does not intrinsically care about their future behavior either in this context.

Let Σ_i^{ls} denote the subset of locally separable strategies and $\varphi_{i,ls}$ the restriction of φ to $\prod_{i \in I} \Sigma_i^{ls}$.

Observation 3 $h \circ \varphi_{i,ls}$ is locally separable whenever h is.

Proof. The observation follows from the fact that $\varphi_{l,s}$ is locally separable (which is easily verified inductively) and the preservation of local separability under composition. \square

Hence, Nash equilibria yield strategy-proof choice functions if the extensive-form mechanism is locally separable and the behavioral strategy-functions are locally separable as functions both of past history and players' types.

The fact that the strong incentive-compatibility implications of the revelation theorem apply to any strategy function in Nash equilibrium, regardless of other co-existing equilibria, suggests that refinements should be of little or no use in a locally separable context. In an extensive-form context, this conjecture is borne out for the notion of subgame-perfection in a straightforward way.

Observation 4 *For any locally separable behavioral strategy-profile s in Nash equilibrium in (h, θ) , there is a history-independent (and thus locally separable) strategy-profile \hat{s} in Nash (and thus: subgame-perfect) equilibrium generating the same play, i.e., such that $\varphi(s) = \varphi(\hat{s})$.*

Proof. This is easily verified for \hat{s} defined by $\hat{s}_i^t(\eta_t) = s_i^t((\varphi(s)_\tau)_{\tau=1, \dots, t-1})$ for all η_t and all t , due to the fact that s_i and \hat{s}_i agree on histories describing individual players' deviations from s . \square

5. EQUIVALENCE OF BAYESIAN NASH AND NASH EQUILIBRIA

Let I now be a countably infinite set. Each agent is endowed with measurable spaces of types (θ_i, Θ_i) , strategies $(\Sigma_i, \mathcal{S}_i)$ and outcome (X_i, \mathcal{X}_i) . Let (Θ, \mathcal{T}) , (Σ, \mathcal{S}) and (X, \mathcal{X}) denote the respective product-spaces; also, let \mathcal{T}^* and \mathcal{S}^* denote the sub-algebra of "tail sets" of \mathcal{T} and \mathcal{S} . Note that mixed strategies can be accommodated in this framework if one includes the conditioning event (the outcome of a privately observed coin-toss, for instance) in the definition of a player's type.

Ex-ante, types are independently distributed according to the product measure $\nu = \prod_{i \in I} \nu_i$, where ν_i is a probability-measure on $(\theta_i, \mathcal{T}_i)$. Ex-interim, players know their type, and their type only. Thus, i 's belief over others' types are given by $\nu_{-i} = \prod_{j \in I \setminus i} \nu_j$.

To compare Bayesian Nash and Nash equilibria, strategic behavior needs to be described by the same type of formal objects in both contexts: strategy-functions that map type profiles to strategy choices. In an incomplete information context, strategy choices may depend only on a player's type: $\sigma_i : \Theta \rightarrow \Sigma_i$ is **incomplete information admissible (i.i.a.)** if $\theta_i = \theta'_i$ implies $\sigma_i(\theta) = \sigma_i(\theta')$, for all $\theta, \theta' \in \Theta$.

For the subsequent results, local separability needs to be strengthened a bit. Rather than merely requiring the irrelevance of microscopic information, i.e., that player i 's outcome not depend on j 's strategy, it will now be demanded that player i 's outcome depend on some well-specified "aggregate" of others' strategies such as the empirical distribution of strategies used. Mathematically, such an aggregate will be represented by a variable in a Polish (i.e. complete separable metric) space.

Thus, the mechanism h is "**regularly locally separable**" ("r.l.s.") if there exists a measurable Polish space (Y, \mathcal{Y}) and measurable functions $\bar{h}_i : \Sigma_i \times Y \rightarrow X_i$ and $g : \Sigma \rightarrow Y$ such that $h_i(s) = \bar{h}_i(s_i, g(s))$ for all $s \in \Sigma$ and such that g is \mathcal{S}^* -

measurable. Note that it is without loss of generality to assume all players' outcome to depend on others' strategies via the same aggregating function g , since the countable product of Polish spaces is Polish (Cohn (1980), p.253).

The strategy-function a is **measurably locally separable** if it is \mathcal{T}^* -measurable – or, equivalently, if it is locally separable as well as \mathcal{T} -measurable.

The crucial mathematical fact is the following:

Lemma 2 *If g is a \mathcal{S}^* -measurable map into the Polish space (Y, \mathcal{Y}) and a is measurably locally separable, then, for some $y \in Y : (g \circ \sigma)(\theta) = y$ ν -almost surely.*

Proof. It is easily seen that $\eta = g \circ a$ is \mathcal{T} -measurable. Thus, as a consequence of Kolmogorov's zero-one-law (Bauer (1981, p.153), for all $V \in \mathcal{Y}_i : (\nu \circ \eta^{-1})(V)$ must be zero or one.

By a well-known implication of the regularity of measures on Polish spaces (Bauer (1981), p.213), this implies that, for some $y \in Y : (\nu \circ \eta^{-1})(\{y\}) = 1$, that is: $\eta(\theta) = y$ ν -almost surely. \square

The twist of the lemma is to combine a context in which Kolmogorov's zero-one property arises, a countably infinite product of measures, with another one, a Polish space, in which this property eliminates uncertainty. Note that, unless ν is degenerate, this "elimination of uncertainty" does not hold for the original context; indeed, there is in general not even a smallest set $T \in \mathcal{T}^*$ of environments such that $\nu(T) = 1$.

Observe also that, in contrast to the symmetric case (see example 2 below), it is obviously not enough to ensure the almost-sure constancy of the (limit-) distribution of strategies.

Finally, note that the validity of the assertion, and thus the restrictiveness of the assumption of regular local separability, hinges on the measurability of g as much as on the Polishness of its range.

Example 1 (Non-measurable g): Assume all Θ_i to be finite and consider any ν such that $\nu(T) = 0$ for all local domains T . Let $a : \Theta \rightarrow \Theta$ be the identity, and $g : \Theta \rightarrow \Theta$ be any function such that

- i) $g(\theta) \sim \theta$ for all $\theta \in \Theta$, and
- ii) $\theta \sim \theta'$ implies $g(\theta) = g(\theta')$ for all $\theta, \theta' \in \Theta$.

(Θ, \mathcal{T}) is Polish as the countable product of finite, hence Polish spaces. Moreover, for all $y \in \Theta$, $(g \circ \sigma)^{-1}(y) = \nu(\{y\})$, which implies $\nu(\{\theta \mid (g \circ \sigma)(\theta) = y\}) = \nu(\{y\}) = 0$. In view of lemma 2, g cannot be $\mathcal{T}^* - \mathcal{Y}$ -measurable. \square

As a consequence of lemma 2, the outcome of a player's strategy-choice is determined essentially deterministically; as a result, a Bayesian Nash equilibrium can be defined in terms of preferences over sure outcomes only. We will write \bar{y} explicitly as a function of ν and \mathbf{a} , $\bar{y} = G(\nu, \mathbf{a})$.

The i.i.a. strategy-function profile \mathbf{a} is in **Bayesian Nash equilibrium (BNE)** in (h, ν) with r.l.s. h if, for no $i \in I$, $\theta \in \Theta_i$ and $s \in \sum_{i'} \bar{h}_i(s, G(\nu, \sigma)) \succ_{\theta} \bar{h}_i(\sigma_i(\theta), G(\nu, \mathbf{a}))$. Note that, as defined, player's preferences over outcomes are determined by their own type. This could be generalized to allow for preferences that depend on other players' types in a locally separable fashion; of course, this would require regularity assumptions on the mapping $\theta \mapsto \succ_{i, \theta}$ based on an appropriate topology on the space of players' preferences as well.

Proposition 2 *If σ is BNE in the game (h, ν) with r.l.s. h , then for some locally closed domain of profiles $\mathcal{D} \in S^*$ with $\nu(\mathcal{D}) = 1$, \mathbf{a} is in NE. Moreover, $h \circ \mathbf{a}$ is strategy-proof on \mathcal{D} .*

Proof. $\mathcal{D} = \{\theta \in \Theta \mid g(\sigma(\theta)) = G(\nu, \sigma)\}$ does the job. Clearly, $\mathcal{D} \in \mathcal{T}^*$ by the S^* -measurability of g and the i.i.a. property of σ .

Moreover, on \mathcal{D} , from its definition and the fact that a is in BNE, for no $\theta \in \mathcal{D}$, $i \in \mathbf{I}$ and $s \in \Sigma_i : h(s, \sigma_{-i}(\theta)) \succ_{\epsilon} h(\sigma(\theta))$. Thus, a is in NE on \mathcal{D} , and $h \circ a$ is strategy-proof on \mathcal{D} . \square

Remark 2: If all strategy-spaces are finite or countably infinite, it suffices to assume the outcome spaces to be Polish to arrive at the conclusion of proposition 2 by a similar argument.

While Bayesian Nash equilibria, in contrast to Nash equilibria, are always characterized by incentive-compatibility constraints, local separability in an infinite context allows to eliminate players' beliefs in their specification. According to proposition 2, this occurs on a domain of "representative" environments \mathcal{D} with ex-ante and ex-interim probability one. The proposition fails to specify, however, what the representative environments are; this is achieved easily if everything is symmetric in players.

Example 2 (Symmetric environment and strategy functions):

Assume all type-, strategy- and outcome-spaces to be identical across players and types to be distributed identically; denote $\Sigma_i = \Sigma_*$, $\Theta_i = \Theta_*$, and $\nu_i = \nu_*$ for $i \in \mathbf{I}$, and assume Σ_* to be Polish.

Let $\mathcal{Y} = \mathcal{M}(\Sigma_*)$ and \mathcal{X} be the Borel σ -algebra obtained from endowing $\mathbf{M}(\mathbb{C},)$ with the weak topology. Identifying \mathbf{I} with the set \mathbf{N} of natural numbers, define $g(s)$ as the limit (in the weak topology) of the empirical frequency distributions of strategies $\{s_1, \dots, s_n\}$ as n tends to infinity, if such a limit exists, and as equal to some constant μ otherwise.

Assume the strategy-functions to be i.i.a. and symmetric, i.e. $\theta_i = \theta'_j$ to imply $\sigma_i(\theta) = \sigma_j(\theta')$, for all $i, j \in \mathbf{I}$, $\theta, \theta' \in \Theta$; then the strategy-function of any player is described by the same function $\sigma_* : \Theta_* \rightarrow \Sigma_*$. By the Glivenko-Cantelli theorem

(see Hildenbrand (1974, p. 52-53)), $(g \circ \sigma)(\theta) = \nu \circ \sigma_*^{-1}$ ν -a.s. . In other words, $G(\nu, \sigma) = \nu \circ \sigma_*^{-1}$. \square

The strategy-function profile σ is "**regularly locally separable**" ("r.l.s.") if there exists a measurable Polish space (Y', \mathcal{Y}') and measurable functions $\bar{\sigma}_i : \Theta_i \times Y' \rightarrow \Sigma_i$ and $f : \Theta \rightarrow Y'$ such that $\sigma_i(\theta) = \bar{\sigma}_i(\theta_i, f(\theta))$ for all $\theta \in \Theta$ and such that f is \mathcal{T}^* -measurable. Define $F(\nu)$ analogously to $G(\nu, a)$ above.

Proposition 2 has the following converse.

Proposition 3 *If $a : \Theta \rightarrow \Sigma$ is r.l.s. and in NE in the r.l.s. mechanism h , and $\bar{\pi}$ types are independently distributed according to ν , there exists a strategy-function profile δ in BNE such that $\sigma(\theta) = \hat{\sigma}(\theta)$ ν -almost surely.*

Proof. Based on the assumed r.l.s. of a , define an i.i.a. strategy-function $\hat{\sigma}_i : \Theta \rightarrow \Sigma_i$ by $\hat{\sigma}_i(\theta) = \bar{\sigma}_i(\theta_i, F(\nu))$ for $\theta \in \Theta$. Since

$$\hat{\sigma}_i(\theta) = \sigma_i(\theta) \text{ whenever } f(\theta) = F(\nu), \quad (1)$$

$\delta_i(\theta) = \sigma_i(\theta)$ ν -almost surely, and thus trivially $G(\nu, \delta) = G(\nu, a)$.

Consider the set \mathcal{D} of environments θ such that

- i) $\hat{\sigma}_i(\theta) = \sigma_i(\theta)$, and
- ii) $g(\sigma(\theta)) = G(\nu, a)$.

\mathcal{D} is non-empty since $\nu(\mathcal{D}) = 1$ in view of (1) and lemma 2.

For a given $i \in \mathbf{I}$ and $\bar{\theta}_i \in \Theta_i$, $i \in \mathbf{I}$, take $\theta \in \mathcal{D}$ such that $\theta_i = \bar{\theta}_i$.

Since a is in NE, for no $s \in \Sigma_i$:

$$\bar{h}_i(s, g(\sigma(\theta))) \succ_{\theta} \bar{h}_i(\sigma_i(\theta), g(\sigma(\theta))),$$

and thus, by the definitional properties i) and ii) of \mathcal{D} , for no $s \in \Sigma_i$:

$$\bar{h}_i(s, G(\nu, \hat{\sigma})) \succ_{\theta_i} \bar{h}_i(\hat{\sigma}_i(\theta), G(\nu, \hat{\sigma})).$$

This shows that the strategy-function profile $\hat{\sigma}$ is in BNE since i and θ_i can be chosen arbitrarily. \square

We have assumed in this section that the set of players is countably infinite. While non-standard, this assumption has sometimes been argued to be conceptually superior to assuming the existence of a continuum of agents. It surely is an easier counterfactual to believe in; Armstrong and Richter (1984) even claim countable models to be more "realistic".

While we have some sympathy with such views, the countability assumption has been made in this paper mainly on grounds of tractability. In view of the problems in invoking laws of large numbers in the continuum uncovered by Judd (1985) and Feldman-Gilles (1985), this is a non-negligible issue. Recently, Al-Najjar (1994) has found a (mathematically sophisticated) way to sidestep these problems in the context of anonymous games; he emphasizes, however, that they remain important in the general case (ibid., p.18).

Finally, it should be noted that under fairly general conditions, Bayesian Nash equilibria have been shown to exist in games of incomplete information with a countable number of players and independent types (Balder-Rustichini (1994)).

REFERENCES

- [1] Al-Najjar, N.I. (1994): "Decomposition and Characterization of Risk with a Continuum of Random Variables," mimeo Université de Québec à Montréal.
- [2] Armstrong, T.E. and M.K. Richter (1984): "The Core-Walras Equivalence," *Journal of Economic Theory* 33, 116-151.
- [3] Balder, E.J. and A. Rustichini (1994): "An Equilibrium Result for Games with Private Information and Infinitely Many Players," *Journal of Economic Theory* 62, 385-393.
- [4] Bauer, H. (1981): *Probability Theory and Elements of Measure Theory*. London. Academic Press.
- [5] Champsaur, P. and G. Laroque (1981): "Fair Allocations in Large Economies," *Journal of Economic Theory* 25, 269-282.
- [6] Cohn, D.C. (1980): *Measure Theory*. Boston, Birkhauser
- [7] Dubey, P., A. Mas-Colell and M. Shubik (1980): "Efficiency Properties of Strategic Market Games: An Axiomatic Approach," *Journal of Economic Theory* 22, 339-362.
- [8] Feldman, M. and C. Gilles (1985): "An Expository Note on Individual Risk without Aggregate Uncertainty," *Journal of Economic Theory* 35, 26-32.
- [9] Gibbard, A.: "Manipulation of Voting Schemes: A General Result," *Econometrica* 41, 587-602.

- [10] Hammond, P. (1979): "Straightforward Individual Incentive Compatibility in Large Economies," *Review of Economic Studies* 46, 263-282.
- [11] Hildenbrand, W. (1974): *Core and Equilibria of a Large Economy*. Princeton, Princeton UP.
- [12] Judd, K.L. (1985): "The Law of Large Numbers with a Continuum of IID Random Variables," *Journal of Economic Theory* 35, 19-25.
- [13] Mas-Colell, A. and X. Vives (1993): "Implementation in Economies with a Continuum of Agents," *Review of Economic Studies* 60, 613-629.
- [14] Maskin, E. (1985): "The Theory of Implementation in Nash Equilibrium: A Survey," in: Hurwicz, L. et al. (eds.): *Social Goals and Social Organization: Essays in the Memory of Elisha Parner*, 173-204, Cambridge, Cambridge UP.
- [15] Nehring, K. (1995): "Incentive-Compatible and Efficient Resource Allocation in Large Economies," mimeo University of California, Davis.
- [16] Roberts, K. (1984): "The Theoretical Limits to Redistribution," *Review of Economic Studies* 47, 587-597.
- [17] Schmeidler, D. (1973): "Equilibrium Points of Non-Atomic Games," *Journal of Statistical Physics* 7, 295-300.