

RANDOM-PLAYER GAMES*

Igal Milchtaich

Department of Economics, Bar-Ilan University, Ramat Gan 52900, Israel

Email: igal.milchtaich@biu.ac.il

Web: <http://faculty.biu.ac.il/~milchti>

Abstract. This paper introduces general games with incomplete information in which the number, as well as the types or identities, of the participating players are determined by chance and might not be known to the players when they make their choices of actions. In these games, the selection of the number and types of players is modeled as a finite point process on a suitable type space. Definitions of pure-strategy, mixed-strategy, and correlated equilibria in random-player games are given, extending the corresponding ones for finite games, Bayesian games, and games with population uncertainty, which may all be considered as special cases of random-player games. *Journal of Economic Literature* Classification Numbers: C72, D80.

Keywords: Incomplete information, point processes, correlated equilibrium, auctions

1. INTRODUCTION

Most models of strategic interactions assume a finite, commonly known number of participating players. Other models involve an infinite number of players. The latter are particularly suitable for the study of interactions involving a very large number of economic agents, each with a negligible effect on the others. A third kind of model involves a finite but *random* number of players. Such models have been employed, for example, in the study of auctions (McAfee and McMillan, 1987; Matthews, 1987; Harstad *et al.*, 1990) and elections (Myerson, 1998a, 1998b, 2000). Among the many other potential applications are the choices of routes in congested networks (e.g., road or computer networks), housing markets, and Internet auctions.

Games with a random number of players are a natural extension of the familiar notion of n -player Bayesian games. Nevertheless, they raise conceptual and modeling issues not encountered in that context. A standard model of an n -player Bayesian game employs a set Ω of possible “states of the world,” a probability measure on this set (the “common prior”) and, for each of the n players, a partition of Ω (the player’s “information partition”). Knowing the state of the world amounts to completely resolving all uncertainty about the payoff function and the information each player has. A player does not generally know the state of the world, but only to which element of his information partition it belongs. Based on this, he constructs his posterior about the actual state of the world as a conditional probability measure on Ω . As the following example shows, this model cannot be easily adapted to the case of a random

* Published in *Games and Economic Behavior* 47 (2004), 353–388. The first version of this paper was written while the author was visiting the Managerial Economics and Decision Sciences Department, J.L. Kellogg Graduate School of Management, Northwestern University.

number of players. Suppose that, with equal probability, there is a single player or one million identical players. Denote by p the probability each player assigns to the existence of other players beside himself. Assuming there is no differential information, the value of p must be the same in all states of the world and for all players. Thus, the posterior probability that there is more than one player is constant, and must therefore be equal to the prior, 0.5. However, a little reflection shows that the answer $p = 0.5$ is wrong. From the perspective of any single player, the odds that he is not alone in the world are one million-to-one. It appears that the reason this standard model fails with a random number of players is that a player can imagine a state of the world in which he *would not exist* as a player in the game. Yet, the description of a player's beliefs by means of an information partition of Ω assigns him beliefs in *all* states of the world. This is not only an absurdity but also, as the above example shows, may lead to incorrect answers.

The transition from a deterministic to a stochastic number of players is particularly problematic if this number is *unbounded* in the sense that it has a positive probability of exceeding any fixed bound, e.g., if it has a geometrical or Poisson distribution. A common-sense approach to modeling such a situation would involve a countably infinite set of potential players, indexed by the natural numbers, and assignment of a particular probability to each finite subset of potential players chosen to play. Arguably, if players are identical in all respects, then they should also all have the same beliefs regarding the number of other players. For these beliefs to be consistent with the way players are actually selected, they should coincide with the players' posteriors. That is, a player's belief that there are n players beside him should equal the conditional probability that the number of players is $n + 1$, given that this particular player was selected to play. However, this paper shows that, irrespective of the probabilities assigned to the various subsets of potential players, with an unbounded number of players, the posterior beliefs are *never* identical. Therefore, this kind of approach inevitably leads to players having non-identical beliefs.

The possibility of a random number of players adds an interesting perspective to the analysis of correlated equilibria in incomplete information games. In *complete* information games, correlated equilibria can be defined in several ways, but the definitions turn out to be essentially equivalent. By contrast, the literature on n -player incomplete information games offers several non-equivalent definitions of correlated equilibrium (see Forges, 1993). Some of these are readily adapted to the general setting of a random number of players, while for others, such as communication equilibrium, the generalization is less straightforward or useful. As this paper shows, in incomplete information games (with either a fixed or stochastic number of players), there is only one definition of correlated equilibrium that includes pure-strategy equilibria as a special case and is preserved under identification of "information types." This means that identification of two or more players with no payoff-relevant differences (who may only differ in what they *know*) transforms each correlated equilibrium in the original game into a correlated equilibrium in the new one. One rationale for requiring this property is that, in complete information games, it uniquely characterizes the class of correlated equilibria.

The main purpose of this paper is to set up a general and coherent framework for the study of games in which the participating players have incomplete information about both the number and types or identities of the other players. The mathematical tool used for modeling the selection of players in such games is finite point processes. A brief account of the relevant

theory is given in Section 2, which also contains a number of original results that are required in the sequel. A finite point process can be thought of as a random selection of a finite subset of some space, possibly with multiple points (i.e., the same space element represented more than once). In the present context, the relevant space (which can be finite or infinite, countable or uncountable) consists of all the possible player types. Two interpretations of player types are possible. According to the first, each type represents a single potential player. According to the second, a type is a list of a player's possible attributes. The set of actual players may include several players with identical attributes. By contrast, any individual potential player can be represented either once or not at all. Hence, the first interpretation constrains the finite point processes selecting the players to be *simple* in the sense that multiple points never occur. A player is assumed to know his own type. However, unless that particular type is always selected with the same number and types of other players, he cannot be certain about the number and types of his companions. The conditional distribution of the number and types of the rest of the players, given a player's type, is that type's posterior. This posterior may also be expressed as a finite point process on the space of player types.

After the players are selected, each of them chooses one of a set of strategies admissible to his type. Without loss of generality, it may be assumed that different player types have disjoint strategy sets. The description of a random-player game is completed by specifying the payoff function, which gives each player's payoff for any strategy he chooses and any combination of strategies for the other players (whose number may vary). The formal definition of a random-player game is given in Section 3, which also includes several examples showing how certain previously introduced models fit into the current setup. In Section 4, several notions of strategy profiles in random-player games are defined and their interrelationships examined. The most general of these is correlated strategy, which allows the players' strategies to depend in an arbitrary manner on the number and types of all the players. For each notion of strategy profile, there is a corresponding notion of equilibrium. In Section 5, pure-strategy, mixed-strategy, and correlated equilibria in random-player games are defined, and three corresponding equilibrium-existence results are described. The sufficient conditions for the existence of mixed-strategy equilibrium in random-player games generalize those of Milgrom and Weber (1985) for n -player Bayesian games. Alternative definitions of correlated equilibrium in random-player games (and other games with incomplete information) are examined in Section 6. Section 7 extends the framework outlined above by allowing players to hold arbitrary *subjective* beliefs about the number and types of the other players (rather than posterior beliefs, which are derived from a common prior). The main aim is to show that a universal belief space, along the lines of Mertens and Zamir (1985), can be developed for games with a random number of players. The paper concludes with an appendix, containing the proofs of the propositions and theorems in it.

2. FINITE POINT PROCESSES

2.1. Fundamentals

A *Polish space* is a topological space homeomorphic to a complete, separable metric space. Every countable (and, in particular, every finite) discrete topological space and every compact metric one are Polish. A function from, or to, a Polish space \mathcal{T} is referred to in this paper as "measurable" if it is measurable with respect to the σ -algebra of Borel sets in \mathcal{T} . The collection of all measurable functions from \mathcal{T} to the nonnegative ray \mathbb{R}_+ is denoted $\mathfrak{F}(\mathcal{T})$. The

term “measurable set”, in this paper, is synonymous with “Borel set”. “Measure” always means a nonnegative σ -additive Borel measure. The *weak topology* on the set of all finite measures on a Polish space \mathcal{T} is the smallest topology on this set with respect to which the function $\mu \mapsto \int f(t) \mu(dt)$ is continuous, for all bounded continuous functions $f: \mathcal{T} \rightarrow \mathbb{R}$. The space $\Delta(\mathcal{T})$ of all probability measures on \mathcal{T} is Polish in the weak topology (Parthasarathy, 1967, Theorems II.6.2 and II.6.5). A function $x \mapsto \mu_x$ from a measurable space \mathcal{X} to $\Delta(\mathcal{T})$ is measurable if and only if the real-valued function $x \mapsto \mu_x(B)$ is measurable for all measurable sets B in \mathcal{T} .

An integer-valued finite measure T on a Polish space \mathcal{T} is called a *finite point measure*. For $t \in \mathcal{T}$, a *Dirac measure* (at t) is a finite point measure δ_t such that $\delta_t(B)$ is 1 if $t \in B$ and 0 if $t \notin B$. Every finite point measure T can be written as a finite sum of Dirac measures, $T = \delta_{t_1} + \delta_{t_2} + \dots + \delta_{t_n}$ (Daley and Vere-Jones, 1988, Proposition 7.1.II). Obviously, $n = T(\mathcal{T})$. If t_1, t_2, \dots, t_n are distinct (or if $n = 0$) then T is said to be *simple*. The function sending each simple finite point measure to its support is a one-to-one correspondence between the set of all simple finite point measures on \mathcal{T} and the collection of all finite subsets of \mathcal{T} . A finite point measure that is not simple can be viewed as a finite subset of \mathcal{T} with *multiple points*.

The space $\mathcal{N}_{\mathcal{T}}$ of all finite point measures on a Polish space \mathcal{T} is Polish in the weak topology (Daley and Vere-Jones, 1988, Corollary 7.1.IV). A sequence T_1, T_2, \dots of finite point measures converges to a limit $T = \delta_{t_1} + \delta_{t_2} + \dots + \delta_{t_n}$ if and only if, for all $1 \leq i \leq n$, there is a sequence t_{i1}, t_{i2}, \dots in \mathcal{T} converging to t_i such that, for all k large enough, $T_k = \delta_{t_{1k}} + \delta_{t_{2k}} + \dots + \delta_{t_{nk}}$ (see Resnick, 1987, p. 144). It follows that $\mathcal{N}_{\mathcal{T}}$ is locally compact if and only if \mathcal{T} is locally compact. Also, for every compact set K in \mathcal{T} and every integer n , the set of all finite point measures T on \mathcal{T} with $T(K) < n$ and $T(\mathcal{T} \setminus K) = 0$ is compact (Daley and Vere-Jones, 1988, Corollary A2.6.V).

A random finite point measure \mathbf{T} on a Polish space \mathcal{T} , i.e., a measurable map from some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to $\mathcal{N}_{\mathcal{T}}$, is called a *finite point process*.¹ The condition that \mathbf{T} is measurable is equivalent to the condition that $\mathbf{T}(B)$ is measurable for all measurable sets B in \mathcal{T} (Daley and Vere-Jones, 1988, Proposition 7.1.VIII). The number of points that a finite point process \mathbf{T} selects in each measurable set B , with each point counted with its multiplicity, is thus a random variable. A finite point process is said to be *simple* if it is almost surely simple-valued. The *distribution* of a finite point process \mathbf{T} is the probability measure $\mathbb{P}\mathbf{T}^{-1}$ on $\mathcal{N}_{\mathcal{T}}$ defined by $\mathbb{P}\mathbf{T}^{-1}(A) = \mathbb{P}(\mathbf{T} \in A)$ (A a measurable subset of $\mathcal{N}_{\mathcal{T}}$). The distribution of \mathbf{T} determines the joint distribution of $\mathbf{T}(B_1), \mathbf{T}(B_2), \dots, \mathbf{T}(B_m)$, for all partitions of \mathcal{T} into a finite number of measurable sets B_1, B_2, \dots, B_m . Conversely, the collection of all such joint distributions completely determines the distribution of \mathbf{T} (Daley and Vere-Jones, 1988, Corollary 6.2.IV). If \mathbf{T} is simple, then its distribution is, in fact, completely determined by the probabilities $\mathbb{P}(\mathbf{T}(B) = 0)$, where B ranges over all measurable sets in \mathcal{T} (Daley and Vere-

¹ Throughout this paper, a random element of a Polish space is denoted by a bold letter. Although such a random element is formally a function, defined on some probability space, its argument will always be suppressed.

Jones, 1988, Theorem 7.3.II). Two finite point processes on a Polish space \mathcal{T} that have the same distribution may be identified with one another, as well as with their common distribution. In this way, the set of all finite point processes on \mathcal{T} is identified with the space $\Delta(\mathcal{N}_{\mathcal{T}})$ of all probability measures on $\mathcal{N}_{\mathcal{T}}$. Thus, in this paper, an equality sign between two finite point processes on the same space indicates that they are *equal in distribution*. A sequence $\mathbf{T}_1, \mathbf{T}_2, \dots$ of finite point processes on \mathcal{T} is said to *converge in distribution* to a finite point process \mathbf{T} if the corresponding sequence of distributions converges in $\Delta(\mathcal{N}_{\mathcal{T}})$ to the distribution of \mathbf{T} . This is equivalent to each of the following two conditions:

- (i) $\int f(t) \mathbf{T}_k(dt) \rightarrow \int f(t) \mathbf{T}(dt)$ in distribution for all bounded continuous functions $f: \mathcal{T} \rightarrow \mathbb{R}$, and
- (ii) $\mathbb{E}[e^{-\int f(t) \mathbf{T}_k(dt)}] \rightarrow \mathbb{E}[e^{-\int f(t) \mathbf{T}(dt)}]$ for all such functions f ,

where \mathbb{E} denotes expectation (Daley and Vere-Jones, 1988, Proposition 9.1.VII). As a special case (namely, $\mathbf{T}_1 = \mathbf{T}_2 = \dots$), each of these criteria gives a necessary and sufficient condition for two finite point processes on the same Polish space to be equal in distribution.

For every $g \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})$, the extended real valued function on $\Delta(\mathcal{N}_{\mathcal{T}})$ (well) defined by $\mathbf{T} \mapsto \mathbb{E}[g(\mathbf{T})]$ (\mathbf{T} a finite point process on \mathcal{T}) is measurable. If $\{\mathbf{T}_x\}_{x \in \mathcal{X}}$ is a family of finite point processes on \mathcal{T} , indexed in some measurable space \mathcal{X} , the function $x \mapsto \mathbf{T}_x$ from \mathcal{X} to $\Delta(\mathcal{N}_{\mathcal{T}})$ is measurable if and only if the function $x \mapsto \mathbb{E}[e^{-\int f(t) \mathbf{T}_x(dt)}]$ is measurable for all bounded continuous functions $f: \mathcal{T} \rightarrow \mathbb{R}_+$ (see Kallenberg, 1986, p. 14). In this case, for each probability measure μ on \mathcal{X} , there is a corresponds finite point process \mathbf{T} on \mathcal{T} , which is uniquely defined by the condition $\mathbb{E}[g(\mathbf{T})] = \int \mathbb{E}[g(\mathbf{T}_x)] \mu(dx)$ for all $g \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})$. The finite point process \mathbf{T} is said to be the *mixture* of \mathbf{T}_x with respect to μ .

A *sample* (or empirical) *process* on a Polish space \mathcal{T} is a finite point process obtained by independently drawing (with replacement) a fixed number n (the *sample size*) of elements of \mathcal{T} according to a common probability measure λ on this space. Thus, a sample process can be written as $\delta_{\mathbf{t}_1} + \delta_{\mathbf{t}_2} + \dots + \delta_{\mathbf{t}_n}$, where $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_n$ are independent and identically distributed random elements of \mathcal{T} . A *mixed sample process* is a generalization of this, in which the sample size is not fixed but a *random variable*, which is independent of $\mathbf{t}_1, \mathbf{t}_2, \dots$ (which are themselves i.i.d. random elements of \mathcal{T} with a common distribution λ). A mixed sample process is simple if and only if the sample size is less than two (i.e., zero or one) almost surely or λ is nonatomic. A mixed sample process in which the sample size has a Poisson distribution is called a *finite Poisson process*. A finite Poisson process can equivalently be defined as a finite point process \mathbf{T} such that, for all measurable sets B in \mathcal{T} , the random variable $\mathbf{T}(B)$ has a Poisson distribution (see Daley and Vere-Jones, 1988, Theorem 2.4.III). A finite point process \mathbf{T} is a finite Poisson process if and only if (i) $\mathbf{T}(\mathcal{T})$ has a Poisson distribution and (ii) $\mathbf{T}(B_1), \mathbf{T}(B_2), \dots, \mathbf{T}(B_m)$ are independent for all finite collections of disjoint measurable sets B_1, B_2, \dots, B_m (Kallenberg, 1986, Exercise 7.6).

A finite point process \mathbf{T} that satisfies condition (ii) above is said to have *independent increments*. A finite point process \mathbf{T} satisfying $\mathbf{T}(\{t\}) = 0$ almost surely for all $t \in \mathcal{T}$ is a finite Poisson process if and only if it is simple and has independent increments (Kallenberg, 1986, Corollary 7.4).

A note on terminology is in order here. This paper makes no use of general (i.e., not necessarily finite) point processes, which may select countably infinite sets of points. However, the definitions of point measure, point process, and Poisson process in all the papers cited here reduce, respectively, to the definitions of finite point measure, finite point process and finite Poisson process as given in this paper in the special case in which \mathcal{T} is compact. Hence, for point processes on compact metrizable spaces, the modifier “finite” can be omitted without risk of confusion.

2.2. Posteriors

For every finite point process \mathbf{T} on a Polish space \mathcal{T} , the extended real valued function $\mathbb{E}\mathbf{T}$ on the measurable sets in \mathcal{T} given by $\mathbb{E}\mathbf{T}(B) = \mathbb{E}[\mathbf{T}(B)]$ is a measure on \mathcal{T} (Daley and Vere-Jones, 1988, p. 188), called the *mean measure* of \mathbf{T} . If the mean measure is finite, then there is a family $\{\mathbf{T}_t\}_{t \in \mathcal{T}}$ of finite point processes on \mathcal{T} such that the function $t \mapsto \mathbf{T}_t$ from \mathcal{T} to $\Delta(\mathcal{N}_{\mathcal{T}})$ is measurable and, for all measurable sets B in \mathcal{T} and all measurable sets A in $\mathcal{N}_{\mathcal{T}}$,

$$(1) \quad \int_B \mathbb{P}(\mathbf{T}_t \in A) \mathbb{E}\mathbf{T}(dt) = \mathbb{E}\left[\int_B 1_{\mathbf{T}-\delta_t \in A} \mathbf{T}(dt)\right].$$

(For a proposition p , 1_p is 1 or 0, according to whether p is true or false, respectively.) Equivalently (see Kallenberg, 1986, p. 84),

$$(2) \quad \int \mathbb{E}[h(t, \mathbf{T}_t)] \mathbb{E}\mathbf{T}(dt) = \mathbb{E}\left[\int h(t, \mathbf{T} - \delta_t) \mathbf{T}(dt)\right] \quad (h \in \mathfrak{F}(\mathcal{T} \times \mathcal{N}_{\mathcal{T}})).$$

These finite point processes are essentially unique. More precisely, if $\{\mathbf{T}'_t\}_{t \in \mathcal{T}}$ is another family of finite point processes as above, then $\mathbf{T}'_t = \mathbf{T}_t$ for $\mathbb{E}\mathbf{T}$ -almost all t . The distribution of the finite point process \mathbf{T}_t can be interpreted as the conditional distribution of $\mathbf{T} - \delta_t$, given that a randomly selected term in the presentation of \mathbf{T} as a sum of Dirac measures turned out to be δ_t . In particular, if \mathbf{T} is simple, then \mathbf{T}_t represents the posterior on the (random) set of *other* points, given that t was selected.

A finite point process \mathbf{T} with finite mean measure will be said to have a *common posterior* if \mathbf{T}_t is the same for $\mathbb{E}\mathbf{T}$ -almost all t . If, moreover, $\mathbf{T}_t = \mathbf{T}$ for $\mathbb{E}\mathbf{T}$ -almost all t , then the common posterior will be said to *coincide with the prior*. Every mixed sample process \mathbf{T} , in which the sample size has a finite, nonzero expectation \bar{n} ($= \mathbb{E}\mathbf{T}(\mathcal{T})$) and the elements in the sample are drawn according to a probability measure $\lambda \in \Delta(\mathcal{T})$, has a common posterior. More specifically, the next theorem shows that, in this case, for λ -almost all t in \mathcal{T} (note that $\lambda = (1/\bar{n}) \mathbb{E}\mathbf{T}$), \mathbf{T}_t is also a mixed sample process, the elements in the sample are also drawn according to λ , and the sample-size distribution is given by

$$\mathbb{P}(\mathbf{T}_t(\mathcal{T}) = n - 1) = (n/\bar{n}) \mathbb{P}(\mathbf{T}(\mathcal{T}) = n) \quad (n = 1, 2, \dots).$$

Thus, the posterior probability of $n - 1$ more elements in the sample is proportional to n times the prior probability that the sample size is n . Therefore, $\mathbf{T}_t = \mathbf{T}$ if and only if $\bar{n} \mathbb{P}(\mathbf{T}(\mathcal{T}) = n - 1) = n \mathbb{P}(\mathbf{T}(\mathcal{T}) = n)$ for all n , which holds if and only if $\mathbb{P}(\mathbf{T}(\mathcal{T}) = n) = (\bar{n}^n/n!) e^{-\bar{n}}$ for all n , that is, if and only if \mathbf{T} is a finite Poisson process. The following theorem asserts that mixed sample processes and finite Poisson processes are, in fact, the *only* finite

point processes for which a common posterior exists and coincides with the prior, respectively.

THEOREM 1 (Kallenberg, 1986, Theorem 11.5 and Exercise 11.1). *A finite point process \mathbf{T} with finite mean measure on a Polish space \mathcal{T} has a common posterior if and only if it is a mixed sample process. In this case, for $\mathbb{E}\mathbf{T}$ -almost all t in \mathcal{T} , \mathbf{T}_t is also a mixed sample process, the distribution of the sample size $\mathbf{T}_t(\mathcal{T})$ satisfies*

$$(3) \quad \mathbb{E}\mathbf{T}(\mathcal{T}) \mathbb{P}(\mathbf{T}_t(\mathcal{T}) = n - 1) = n \mathbb{P}(\mathbf{T}(\mathcal{T}) = n) \quad (n = 1, 2, \dots),$$

and, if $\mathbb{E}\mathbf{T}(\mathcal{T}) \neq 0$, each element in the sample is drawn from \mathcal{T} according to the probability measure $(1/\mathbb{E}\mathbf{T}(\mathcal{T})) \mathbb{E}\mathbf{T}$. The common posterior coincides with the prior if and only if \mathbf{T} is a finite Poisson process.

A weaker property than the existence of a common posterior is a common posterior on the total *number of points*. This means that the distribution of the random variable $\mathbf{T}_t(\mathcal{T})$ (which expresses the posterior on the total number of points minus one) is the same for $\mathbb{E}\mathbf{T}$ -almost all t . A finite point process with this property need not be a mixed sample process. However, as the next proposition shows, it shares certain aspects of such point processes. A random variable will be said to be *unbounded* if it has positive probability of exceeding any fixed integer n .

PROPOSITION 1. *Let \mathbf{T} be a finite point process with finite mean measure on a Polish space \mathcal{T} . If the distribution of $\mathbf{T}_t(\mathcal{T})$ is the same for $\mathbb{E}\mathbf{T}$ -almost all t in \mathcal{T} , then:*

- (i) *Eq. (3) holds;*
- (ii) *if $\mathbf{T}(\mathcal{T})$ (or, equivalently, $\mathbf{T}_t(\mathcal{T})$) is unbounded then, for every measurable set B , $\mathbf{T}(B)$ is unbounded or is almost surely equal to zero; and,*
- (iii) *if \mathbf{T} is simple and $\mathbf{T}(\mathcal{T})$ is unbounded, then $\mathbb{E}\mathbf{T}$ is nonatomic.*

If the distribution of $\mathbf{T}_t(\mathcal{T})$ is equal to that of $\mathbf{T}(\mathcal{T})$ for $\mathbb{E}\mathbf{T}$ -almost all t , then $\mathbf{T}(\mathcal{T})$ has a Poisson distribution, and $\mathbb{E}\mathbf{T}$ is nonatomic or \mathbf{T} is not simple.

2.3. Construction of point processes through transition probabilities

The construction of mixed sample processes can be considerably generalized by allowing the selection of points according to different distributions. Specifically, let \mathcal{S} and \mathcal{T} be Polish spaces and, for each $t \in \mathcal{T}$, let μ_t be a probability measure on \mathcal{S} , such that the function $t \mapsto \mu_t$ from \mathcal{T} to $\Delta(\mathcal{S})$ is measurable (or, equivalently, the real-valued function $t \mapsto \mu_t(B)$ is measurable for all measurable sets B in \mathcal{S}). Such a function $t \mapsto \mu_t$ is called a *transition probability*. For every finite point measure $T = \delta_{t_1} + \delta_{t_2} + \dots + \delta_{t_n}$ on \mathcal{T} , let s_1, s_2, \dots, s_n be independent random elements of \mathcal{S} with the joint distribution $\mu_{t_1} \times \mu_{t_2} \times \dots \times \mu_{t_n}$, and let S_T denote the finite point process $\delta_{s_1} + \delta_{s_2} + \dots + \delta_{s_n}$. It is easy to show, using the criterion given in Section 2.1, that the function $T \mapsto S_T$ from $\mathcal{N}_{\mathcal{T}}$ to $\Delta(\mathcal{N}_{\mathcal{S}})$ is measurable. Therefore, for

every finite point process \mathbf{T} on \mathcal{T} , \mathbf{S}_T can be mixed with respect to the distribution of \mathbf{T} , yielding a finite point process \mathbf{S} on \mathcal{S} . By definition, the distribution of \mathbf{S} is given by

$$(4) \quad \mathbb{E}[g(\mathbf{S})] = \mathbb{E}_{\mathbf{T}} [\mathbb{E}_{\mathbf{S}_T} [g(\mathbf{S}_T)]] \quad (g \in \mathfrak{F}(\mathcal{N}_{\mathcal{S}})),$$

and it follows that the mean measure of \mathbf{S} satisfies

$$(5) \quad \int f(s) \mathbb{E}\mathbf{S}(ds) = \int [\int f(s) \mu_t(ds)] \mathbb{E}\mathbf{T}(dt) \quad (f \in \mathfrak{F}(\mathcal{S})).$$

For any fixed transition probability $t \mapsto \mu_t$, the function $\mathbf{T} \mapsto \mathbf{S}$ from $\Delta(\mathcal{N}_{\mathcal{T}})$ to $\Delta(\mathcal{N}_{\mathcal{S}})$, which sends each finite point process \mathbf{T} on \mathcal{T} to the finite point process \mathbf{S} on \mathcal{S} defined by (4), is measurable. If \mathbf{T} is a mixed sample process, then so is \mathbf{S} . In this case, both processes have the same sample-size distribution, and if \mathbf{T} is obtained by sampling elements of \mathcal{T} according to a probability measure λ , \mathbf{S} is obtained by sampling elements of \mathcal{S} according to the probability measure μ defined by $\mu(B) = \int \mu_t(B) \lambda(dt)$. In particular, if \mathbf{T} is a finite Poisson process, then so is \mathbf{S} . Hence, in this case, \mathbf{S} has independent increments and, for all measurable sets B in \mathcal{S} , $\mathbf{S}(B)$ has a Poisson distribution with parameter $\mathbb{E}\mathbf{T}(\mathcal{T}) \int \mu_t(B) \lambda(dt)$.

The following theorem is useful for establishing that a finite point process \mathbf{S} on a Polish space \mathcal{S} can be obtained from a finite point process on another Polish space \mathcal{T} through some transition probability $t \mapsto \mu_t$, without actually having to identify that transition probability. In the theorem, $\{\mathbf{S}_s\}_{s \in \mathcal{S}}$ is (a version of) the family of posteriors, defined as in Eq. (1), with \mathbf{S} and s replacing \mathbf{T} and t , respectively.

THEOREM 2. *Let \mathbf{S} be a finite point process with finite mean measure on a Polish space \mathcal{S} , and τ a measurable function from \mathcal{S} onto another Polish space \mathcal{T} . If for every t in \mathcal{T} there is a finite point process \mathbf{S}_t on \mathcal{S} such that, for every s ,*

$$(6) \quad \mathbf{S}_s = \mathbf{S}_{\tau(s)},$$

then there is a finite point process \mathbf{T} on \mathcal{T} such that \mathbf{S} is obtained from \mathbf{T} through some transition probability $t \mapsto \mu_t$ satisfying

$$(7) \quad \mu_t(\tau^{-1}(\{t\})) = 1 \text{ for } \mathbb{E}\mathbf{T}\text{-almost all } t.$$

In the special case in which \mathcal{T} is a one-element set, Theorem 2 reduces to the result that a common posterior implies that \mathbf{S} is a mixed sample process (see Theorem 1). It may, therefore, be viewed as a generalization of that result. The next proposition extends the previous one by establishing the converse of the assertion in Theorem 2 and providing more details about the finite point processes involved.

PROPOSITION 2. *Let τ be a measurable function from a Polish space \mathcal{S} onto another Polish space \mathcal{T} , and \mathbf{S} a finite point process with finite mean measure on \mathcal{S} that is obtained from a finite point process \mathbf{T} on \mathcal{T} through a transition probability $t \mapsto \mu_t$ satisfying (7). For every $t \in \mathcal{T}$, let \mathbf{S}_t be the finite point process on \mathcal{S} obtained from \mathbf{T}_t through that transition probability. Then:*

- (i) Eq. (6) holds for $\mathbb{E}\mathcal{S}$ -almost all s ;
- (ii) $T = S \circ \tau^{-1}$; and
- (iii) T is a mixed sample process if and only if S is a mixed sample process, and in this case, both processes have the same sample-size distribution.

A finite point process S on a Polish space \mathcal{S} that can be obtained from some finite point process T on another Polish space \mathcal{T} through a transition probability $t \mapsto \mu_t$ with *countable range* (i.e., only a finite or countably infinite number of distinct μ_t 's) will be called an *extended sample process*. Every mixed sample process is also an extended sample process (with \mathcal{T} a one-element set). Also, every finite point process S on a countable Polish space \mathcal{S} is an extended sample process (with $\mathcal{T} = \mathcal{S}$, $T = S$, and $\mu_t = \delta_t$ for all t). It is not difficult to see that the *sum* of any finite number of independent extended sample processes on the same space is also an extended sample process. In addition, if S_1, S_2, \dots is any finite or infinite list of extended sample processes on the same Polish space \mathcal{S} , any *mixture* of these processes (see Section 2.1) is also an extended sample process. In fact, *any* extended sample process can be presented as such a mixture, in which each S_k is the sum of a finite number of independent sample processes, all with a sample size of one.

3. DEFINITION AND EXAMPLES OF RANDOM-PLAYER GAMES

The space of *strategies* is a Polish space \mathcal{S} . A *play* is a finite point measure on \mathcal{S} . The *payoff function* is a bounded and measurable function $u : \mathcal{S} \times \mathcal{N}_{\mathcal{S}} \rightarrow \mathbb{R}$, which specifies the payoff $u(s, S)$ of a player who uses strategy s when the other players' play is S . Implicit in this definition is the assumption that a strategy conveys all relevant information about the player using it. This assumption involves little loss of generality, since the strategy sets of different types of players can always be made disjoint by tagging their elements differently. The association of strategies with player types is formally given by a *type map*, which is a retraction $\tau : \mathcal{S} \rightarrow \mathcal{T}$ (i.e., a continuous function satisfying $\tau(\tau(s)) = \tau(s)$ for all s) with the property that $\tau^{-1}(K)$ is compact for every compact set $K \subseteq \mathcal{S}$.² The *type space*, denoted by \mathcal{T} , is the range of this map (i.e., $\tau(\mathcal{S})$). For each player type $t \in \mathcal{T}$, the *strategy set* of t is the compact set $\tau^{-1}(\{t\})$. The interpretation is that each player type t is identified with one of that type's strategies, to which all the other strategies are mapped by τ . It makes little difference *which* of t 's strategies is chosen for this purpose. Indeed, if $\tau' : \mathcal{S} \rightarrow \mathcal{S}$ is another retraction such that, for all s and s' , $\tau(s) = \tau(s')$ if and only if $\tau'(s) = \tau'(s')$, then \mathcal{T} and $\tau'(\mathcal{S})$ are homeomorphic (and the restriction of τ' to \mathcal{T} is a homeomorphism between them). Being a retract, the type space \mathcal{T} is a closed subset of \mathcal{S} , and is therefore Polish in the relative topology.

² This assumption implies that the type map is closed. Equivalently, the multivalued function (or correspondence) $t \mapsto \tau^{-1}(\{t\})$, defined on the range of τ (which is denoted below by \mathcal{T}), is upper semicontinuous. To see that, consider a sequence t_1, t_2, \dots in \mathcal{T} converging to some limit t . Since the closure K of a converging sequence is compact, so is its inverse image $\tau^{-1}(K)$. Hence, every sequence s_1, s_2, \dots in \mathcal{S} such that $\tau(s_i) = t_i$ for all i has a subsequence converging to some limit s . By continuity of τ , $\tau(s) = t$. This shows that, if $F \subseteq \mathcal{S}$ is a closed set and $t_i \in \tau(F)$ for all i , then $t \in \tau(F)$. Hence, $\tau(F)$ is closed.

A *game* is a finite point measure $T = \delta_{t_1} + \delta_{t_2} + \dots + \delta_{t_n}$ on \mathcal{T} . The elements t_1, t_2, \dots, t_n are the types of players in the game. These types need not be distinct: If T is not simple, then two or more players have the same type. A *random-player game* is a finite point process \mathbf{T} with finite mean measure on \mathcal{T} . The mean measure $\mathbb{E}\mathbf{T}$ gives the expected number of players belonging to each set of types.

A random-player game \mathbf{T} has *independent types* if it is a mixed sample process. As the first part of Theorem 1 shows, this is equivalent to the existence of a common posterior.

COROLLARY 1. *A random-player game has independent types if and only if all players share the same posterior beliefs on the number and types of the other players.*

Example 1. Finite games. Every finite game, and more generally every n -player game in which the strategy sets are compact metric spaces and the payoff functions are bounded and measurable, may be viewed as a random-player game with a degenerate distribution.

Example 2. Bayesian games. Consider an n -player Bayesian game in which, for each player i , the set of possible types is a Polish space \mathcal{T}_i , the set of available actions is a compact metric space \mathcal{A}_i , and the payoff function $u_i: (\times_j \mathcal{T}_j) \times (\times_j \mathcal{A}_j) \rightarrow \mathbb{R}$ is bounded and measurable. The players' (consistent) beliefs are derived from a common prior η on $\times_{i=1}^n \mathcal{T}_i$. To present this as a random-player game, define the space of strategies \mathcal{S} as the disjoint union of $\mathcal{T}_1 \times \mathcal{A}_1, \mathcal{T}_2 \times \mathcal{A}_2, \dots, \mathcal{T}_n \times \mathcal{A}_n$ and the type map τ as the function sending each element (t_i, a_i) of $\mathcal{T}_i \times \mathcal{A}_i$ ($i = 1, 2, \dots, n$) to (t_i, a_i^*) , where a_i^* is some fixed element of \mathcal{A}_i . The type space \mathcal{T} is thus the disjoint union of $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n$. The payoff function is defined as follows: If the play is $\delta_{(t_1, a_1)} + \delta_{(t_2, a_2)} + \dots + \delta_{(t_n, a_n)}$, with $(t_i, a_i) \in \mathcal{T}_i \times \mathcal{A}_i$ for all i , then the payoff $u((t_i, a_i), \sum_{j \neq i} \delta_{(t_j, a_j)})$ of the player whose strategy is (t_i, a_i) is given by $u_i(t_1, t_2, \dots, t_n; a_1, a_2, \dots, a_n)$. For any other strategy and play, $u = 0$ (say). Finally, \mathbf{T} is the simple point process on \mathcal{T} selecting exactly one point t_i in each set \mathcal{T}_i , such that the joint distribution of t_1, t_2, \dots, t_n is η .

Example 3. Games with population uncertainty (Myerson, 1998a). In a game with population uncertainty, there is a finite set of player types \mathcal{T} and a finite set of actions \mathcal{A} . The number of players of each type is a random variable. Each player chooses an action in \mathcal{A} , and his payoff is determined as a bounded function by his type, his action, and the number of other players choosing each action. A random-player game corresponding to this game can be constructed as follows. The space of strategies \mathcal{S} is the Cartesian product $\mathcal{T} \times \mathcal{A}$. The type map τ , as in the previous example, is essentially the projection on the first coordinate, and the type space is thus identifiable with \mathcal{T} . The random vector of the number of players of each type defines a finite point process \mathbf{T} on \mathcal{T} . Assuming that the total number of players has finite expectation, \mathbf{T} has a finite mean measure. The assumption that each player's payoff depends only on the number of other players choosing each action (and not on their types) is expressed by the condition that $u(s, S) = u(s, S')$ whenever S and S' have the same marginal on \mathcal{A} .

A *symmetric game with local interactions* (Mailath *et al.*, 1997) is a special kind of game with population uncertainty. Here, the elements of \mathcal{T} are not distinct types of players but identical individual potential players—candidates for playing some fixed symmetric n -player game G . The identity of the actual players is determined as a simple finite point process \mathbf{T} on \mathcal{T} , in

which exactly n individuals are always selected. For example, the selection of the players may reflect their spatial distribution, with only neighbors playing against each other. Because of the symmetry of G , unlike a general game with population uncertainty, a player's type (i.e., identity) has no effect on his payoff.

Another special kind of game with population uncertainty is one in which the numbers of players of different types are independent, i.e., T has *independent increments* (see Section 2). This does *not* imply that the game has independent types, or vice versa. (The meaning of independent types is that each *player* is assigned a type independently of the others.) Indeed, players of different types may have different posteriors even regarding the *number* of other players (cf. Corollary 1). It can be shown that, in a game of this kind in which at least two player types occur with positive probability, a necessary and sufficient condition for independent types is that the numbers of players of the various types are independent Poisson random variables. In this case, the game is called a *Poisson game* (Myerson, 1998a, 2000). In such a game, not only do the players share a common posterior, but it also coincides with the prior. In other words, each player's beliefs about the number of *other* players and their types coincide with the prior distribution of the *total* number of players and their types. Myerson (1998a) calls this property "environmental equivalence," and shows that, in the class of games with population uncertainty, it holds *only* for Poisson games. The last part of Theorem 1 shows that a similar result holds more generally.

COROLLARY 2. *A random-player game T has environmental equivalence if and only if T is a finite Poisson process.*

If the posterior beliefs of all the players about the *number* of other players are equal to the prior distribution of the total number of players, then, by the last part of Proposition 1, this number must be a Poisson random variable. However, in this case, unlike the previous one in which the posterior beliefs on both the number and *types* of the other players coincides with the prior, T itself does not have to be a finite Poisson process. For example, it may be a mixture of several finite Poisson processes with the same sample-size distribution.

The potential importance of finite Poisson processes follows from the fact (spelled out in a more precise manner in Section 2.1) that any simple finite point process on a Euclidean space, say, in which (i) the numbers of points selected from disjoint regions in the space are independent and (ii) the probability of selecting any given single point is zero is *necessarily* a finite Poisson process. Rare events occurring independently in different places or to different people roughly satisfy these two conditions.

Example 4. *Extended Poisson games.* (The following definition is somewhat broader than given by Myerson, 1998b.) In an extended Poisson game, the number and types of the players are determined as follows. There is a set of player types \mathcal{T} , which is a Polish space (e.g., finite, countably infinite, or compact metric space), a finite set of actions \mathcal{A} , and a finite set of possible states of the world Ω . First, a state of the world ω is selected according to some fixed probability measure on Ω . Then, the number and types of the players are determined as a state-dependent finite Poisson process T_ω on \mathcal{T} . Each player chooses an action, and his payoff is determined as a bounded and continuous function by his type, his action, the number of other players choosing each action, and the state of the world. To present this as a random-player game, let the space of strategies \mathcal{S} be defined as the disjoint union of Ω and $\mathcal{T} \times \mathcal{A}$.

The type map is defined by $\tau(\omega) = \omega$ ($\omega \in \Omega$) and $\tau(t, a) = (t, a^*)$ ($(t, a) \in \mathcal{T} \times \mathcal{A}$), where a^* is some fixed element of \mathcal{A} . The type space is thus the disjoint union of Ω and $\mathcal{T} \times \mathcal{A}$. T is a finite point process on this space selecting, in the way described above, a single state of the world and a combination of player types. From the remarks at the end of Section 2.3, it follows that T is an extended sample process. The payoff function, which is defined in the obvious way, is bounded and continuous and satisfies $u(s, S) = u(s, S')$ whenever $S(\{\omega\}) = S'(\{\omega\})$ for all $\omega \in \Omega$ and $S(\mathcal{T} \times \{a\}) = S'(\mathcal{T} \times \{a\})$ for all $a \in \mathcal{A}$ (i.e., whenever the state of the world and the number of players choosing each action is the same in S and S').

The construction in the last example suggests a general way of modeling the effect on the players' payoffs of random parameters, or states of the world, which may be unknown to any of them. Namely, a fictitious player type, with a singleton strategy set, is associated with each state of the world. Note, however, that such a construction is necessary only if there is a *residual* uncertainty about the payoffs that remains even if the number and types of players are known. In many applications, such residual uncertainty is inconsequential, since whatever the players cannot know *in principle*, even by pooling their knowledge, cannot have any effect on their behavior, and may, therefore, be integrated away by taking expectations. See, however, the discussion of correlated equilibrium in Section 6.

Example 5. *Auctions with a stochastic number of bidders* (McAfee and McMillan, 1987; Matthews, 1987; Harstad *et al.*, 1990). The set of *potential bidders* is finite or countably infinite. It is indexed by (a subset of) the natural numbers $\{1, 2, 3, \dots\}$. The number of *active bidders* is random and finite, with finite expectation. The active bidders are selected from the potential ones in a way that satisfies the following *symmetry assumption*: All potential bidders, when selected, have the same posterior beliefs on the number of other bidders. A bidder's type is a random variable, the value of which lies in some closed interval. Different bidders may have dependent (e.g., affiliated) or independent types. However, an active bidder's type does not give him any information about which other potential bidders are active. A bidder's valuation of the indivisible asset being auctioned is a function of his type only (i.e., private values) or of the types of all the active bidders. Each of the active bidders submits a bid to the auctioneer. Together, these bids determine the bidders' costs and the identity of the winner. First- and second-price sealed-bid auctions are examples of concrete applicable mechanisms.

This model of auctions with a stochastic number of bidders has a limitation, which is not obvious at first sight. As McAfee and McMillan (1987) show, for any *bounded* (i.e., not unbounded) random variable N whose values are nonnegative integers, the active bidders can be selected in such a way that each bidder's posterior on the number of other bidders is given by N . However, it follows from the results in this paper that this cannot be done for *any* unbounded random variable. In other words, if there is no upper bound to the number of active bidders (e.g., their number has a geometric or Poisson distribution), then the symmetry assumption can *never* be satisfied. This can be shown as follows. The selection of active bidders in the setup described above can be viewed as a simple finite point process on the natural numbers. Since this space is countable, the mean measure of the process is purely atomic. Therefore, by (iii) in Proposition 1, the symmetry assumption is inconsistent with an unbounded number of active bidders. With a *continuum* of potential bidders, this limitation on the number of active bidders would not exist. Indeed, for any random variable N whose values are nonnegative integers and any uncountable Polish space \mathcal{T} , there is a simple finite

point process on \mathcal{T} such that each player's posterior on the number of other players is given by the distribution of N . For example, it follows from (3) that, for any nonatomic probability measure λ on \mathcal{T} , the following simple mixed sample process has this property: for every $n \geq 1$, the probability that the sample size is equal to n is $(1/n) \mathbb{P}(N = n - 1)$; and the elements in the sample are drawn independently according to λ .

These remarks may be generalized. In essence, they refer to an alternative formulation of games with a stochastic number of players. This formulation stems from the following interpretation of n -player Bayesian games (Harsanyi, 1967). Initially, the n players, like manikins, lack any specific features. Then, the attributes of the players are randomly selected according to some probability measure on $\mathcal{T}_1 \times \mathcal{T}_2 \times \cdots \times \mathcal{T}_n$, where \mathcal{T}_i is the set of all possible types for player i . An analogous model for games with a random number of players would consist of a countably infinite set of potential players, indexed by the natural numbers, and a probability measure on the product space $\mathcal{T}^*_1 \times \mathcal{T}^*_2 \times \cdots$, where \mathcal{T}^*_i consists of all real types of player i plus a “dummy” type, whose actions do not affect the players' payoffs. Assuming that, with probability one, there is only a finite number of potential players with a non-dummy type, this probability measure simultaneously specifies the random set of players—which consists of all potential players who are of a non-dummy type—and their types. However, this setup is less general than random-player games. In particular, it does not allow players to have an arbitrary common posterior on the number of other players. Specifically, any common posterior must assign zero probability to this number exceeding some fixed upper bound. This follows from Proposition 1: For any probability measure on $\mathcal{T}^*_1 \times \mathcal{T}^*_2 \times \cdots$, there is a corresponding simple finite point process on the disjoint union \mathcal{T} of the players' sets of real types, which always selects either zero or one element from each of these sets. By (ii) in Proposition 1, this implies that, if the number of players is unbounded, they cannot have a common posterior on the number of other players.

4. STRATEGY PROFILES

A measurable function $\sigma: \mathcal{T} \rightarrow \mathcal{S}$ such that $\tau \circ \sigma$ is the identity map on \mathcal{T} is called a *pure-strategy profile*. It assigns each player type t one of that type's strategies. The inclusion map from \mathcal{T} to \mathcal{S} (which maps each t to itself) is an example of a pure-strategy profile, indeed a continuous one. A *mixed strategy* for player type t is any probability measure μ_t on \mathcal{S} that is supported in t 's strategy set (i.e., $\mu_t(\tau^{-1}(\{t\})) = 1$). A *mixed-strategy profile* is any transition probability $t \mapsto \mu_t$ such that, for each player type t , μ_t is a mixed strategy for t . For a given random-player game \mathbf{T} , a *distributional strategy* is any measure μ on \mathcal{S} such that $\mu \circ \tau^{-1}$ is equal to $E\mathbf{T}$, the mean measure of \mathbf{T} . A *correlated strategy* is any finite point process \mathbf{S} on \mathcal{S} that satisfies $\mathbf{S} \circ \tau^{-1} = \mathbf{T}$.³ This definition of correlated strategy is rather broad: it does not put any limitations on which strategies are assigned to each combination of player types. Each player may always be assigned the same strategy, or the strategies may vary in an independent

³ Recall that, if $S = \delta_{s_1} + \delta_{s_2} + \cdots + \delta_{s_n}$ is a finite point measure on \mathcal{S} , then $S \circ \tau^{-1} = \delta_{\tau(s_1)} + \delta_{\tau(s_2)} + \cdots + \delta_{\tau(s_n)}$.

or correlated manner. In short, a correlated strategy is any random *play* consistent with the given random-player game. If T is simple, then it can be shown that every correlated strategy is also simple.

These notions of strategies and profiles are related to one another in the usual ways. Namely,

- (i) pure-strategy profiles may be viewed as a special kind of mixed-strategy profile,
- (ii) distributional strategies are essentially just another way of representing mixed-strategy profiles, and
- (iii) mixed-strategy profiles may be viewed as a special kind of correlated strategy.

More specifically:

- (i) Every pure-strategy profile σ is represented by a unique mixed-strategy profile, namely, $t \mapsto \delta_{\sigma(t)}$; a distributional strategy, $\mu = \mathbb{E}T \circ \sigma^{-1}$; and a correlated strategy, $S = T \circ \sigma^{-1}$.
- (ii) Every mixed-strategy profile $t \mapsto \mu_t$ is represented by a unique distributional strategy μ , which is defined by

$$(8) \quad \int f(s) \mu(ds) = \int [\int f(s) \mu_t(ds)] \mathbb{E}T(dt) \quad (f \in \mathfrak{F}(\mathcal{S})).$$

Conversely, for every distributional strategy μ , there is a mixed-strategy profile $t \mapsto \mu_t$ (namely, a regular conditional probability distribution; see Parthasarathy, 1967, Theorem V.8.1) for which (8) holds. This representation of μ as a mixed-strategy profile is essentially unique: For any mixed-strategy profile $t \mapsto \mu'_t$ that satisfies a similar condition, $\mu'_t = \mu_t$ for $\mathbb{E}T$ -almost all t .

- (iii) Every mixed-strategy profile $t \mapsto \mu_t$ is represented by a unique correlated strategy S , which is the finite point process on \mathcal{S} obtained from T through $t \mapsto \mu_t$ (i.e., the finite point process defined by (4)). In other words, S is the random play that results from first selecting the number and types of the players according to T , and then, independently for each player, selecting an element in the strategy set of that player's type t according to the mixed strategy μ_t . It follows from (ii) that (the distribution of) S is, in fact, completely determined by the distributional strategy μ defined in (8). Conversely, this distributional strategy can easily be determined from S . Indeed, comparing (8) and (5) shows that

$$(9) \quad \mu = \mathbb{E}S.$$

Therefore, the function $\mu \mapsto S$, which sends each distributional strategy to the correlated strategy representing it, is one-to-one. Distributional strategies (or, equivalently, mixed-strategy profiles) may thus be seen as a special kind of correlated strategy.⁴

⁴ However, the relative topology on the set of distributional strategies, when these are seen as correlated strategies, is generally stronger than the weak topology on this set. In other words, the function sending each distributional strategy to the correlated strategy representing it is generally not continuous. Indeed, while this function, as shown by Lemma 4 below, has a compact domain, its range

For a given correlated strategy \mathcal{S} in a random-player game \mathbf{T} , consider the posterior \mathcal{S}_s of a player who uses strategy s on the other players' play. This posterior is defined as in (1), with \mathcal{S} and s replacing \mathbf{T} and t , respectively. If players choose their strategies according to some mixed-strategy profile $t \mapsto \mu_t$ (which is related to the correlated strategy \mathcal{S} as in (iii) above), then, clearly, each player's posterior beliefs are completely determined by his type; a player's realized strategy does not give him any additional information about the others' strategies. More specifically, Proposition 2 shows that, in this case, for a player of type t , the posterior on the other players' play is given by the finite point process \mathcal{S}_t obtained from the posterior \mathbf{T}_t on the other players' number and types through the mixed-strategy profile. In other words, each player constructs his posterior beliefs about the other players' play from his beliefs about their number and types, by assuming that each of them independently chooses an action according to the mixed strategy of that player's type. This, of course, is hardly surprising. However, Theorem 2 shows that the *converse* is also true. Specifically, together with Proposition 2, it gives the following corollary.

COROLLARY 3. *The players' posterior beliefs on the other players' play are completely determined by their own types if and only if the players choose their strategies independently according to some mixed-strategy profile.*

It follows from Corollary 3 that, for a correlated strategy \mathcal{S} that does *not* represent a mixed-strategy profile, there is always a strategy s such that the posterior \mathcal{S}_s of a player using s on the other players' play differs from that of a player of *the same type* who uses some other strategy s' . Thus, for a player of this type, his strategy is not independent of the others' play. In fact, it may not even be independent of their number or types. (Recall that the players' strategies identify their types.)

In a random-player game \mathbf{T} with *independent types*, if the players choose their strategies according to some mixed-strategy profile, then they all have the same posterior beliefs about the other players' play. An interesting special case occurs if \mathbf{T} is a finite Poisson process. Here, the common posterior coincides with the prior. That is, each player's beliefs about the *other* players' play are the same as those of an outside observer about *all* the players' play. This follows immediately from (iii) in Proposition 2, which implies that, for a game of this kind, any correlated strategy \mathcal{S} that represents a mixed-strategy profile is a finite Poisson process. Since finite Poisson processes have independent increments (see Section 2), this also proves the "*independent-action property*" of Poisson games (Myerson, 1998a): When the players in such a game choose their actions independently according to some type-dependent mixed strategies, the number choosing each action is a Poisson random variable, and these random variables are independent.

may not even be closed in $\Delta(\mathcal{N}_s)$. It can be shown (cf. the proof of Theorem 3) that a sufficient condition for this function to be continuous is that the distribution of \mathbf{T} is absolutely continuous with respect to the distribution of some extended sample process.

5. EQUILIBRIUM

A correlated strategy S will be said to be a *correlated equilibrium* if $\mathbb{E}[u(s, S_s)] = \sup_{s' \in \tau^{-1}(\{\tau(s)\})} \mathbb{E}[u(s', S_s)]$ for $\mathbb{E}S$ -almost all s . In other words, in a correlated equilibrium, (almost) every strategy s used is a best response for the player using it against his posterior on the other players' play. For a discussion of this notion of correlated equilibrium in random-player games, as well as several alternative ones, see the next section. A mixed-strategy profile is a *mixed-strategy equilibrium* if the correlated strategy S representing it is a correlated equilibrium. An equivalent definition is the following: A mixed-strategy profile $t \mapsto \mu_t$ is a mixed-strategy equilibrium if, for $\mathbb{E}T$ -almost all player types t , μ_t is supported in $\operatorname{argmax}_{s \in \tau^{-1}(\{\tau(t)\})} \mathbb{E}[u(s, S_t)]$. Thus, in a mixed-strategy equilibrium, the mixed strategy of (almost) every player type is a best response against this type's posterior on the other players' play. A pure-strategy profile σ is a *pure-strategy equilibrium* if the corresponding correlated strategy $S = T \circ \sigma^{-1}$ is a correlated equilibrium or, equivalently, if $\mathbb{E}[u(\sigma(t), T_t \circ \sigma^{-1})] = \sup_{s \in \tau^{-1}(\{\tau(t)\})} \mathbb{E}[u(s, T_t \circ \sigma^{-1})]$ for $\mathbb{E}T$ -almost all t .

Milgrom and Weber (1985) showed that, for an n -player Bayesian game satisfying the assumptions in Example 2, a sufficient condition for the existence of mixed-strategy equilibrium is that (i) each player's payoff function u_i is continuous and (ii) the common prior η is absolutely continuous with respect to some product measure on $\times_{i=1}^n \mathcal{T}_i$. A similar result holds for general random-player games T .

THEOREM 3. *Suppose that the payoff function u is continuous. If the distribution of T is absolutely continuous with respect to the distribution of some extended sample process, then a mixed-strategy equilibrium exists.*

Every random-player game T with independent types satisfies the absolute continuity condition in Theorem 3. The same is true for every random-player game T in which the mean measure $\mathbb{E}T$ is supported in a countable subset of \mathcal{T} . Indeed, in this case, T itself is an extended sample process. It follows that if the type space \mathcal{T} is countable, then a sufficient condition for the existence of mixed-strategy equilibrium is that the payoff function is continuous. In the special case in which the type space is countable and each player type has a finite strategy set, it is possible to dispense with the continuity condition. Indeed, as the proof of the following corollary of Theorem 3 shows, the payoff function, in this special case, may be *assumed* to be continuous.

PROPOSITION 3. *Every random-player game with a finite or countably infinite type space and finite strategy sets has a mixed-strategy equilibrium.*

Sufficient conditions for the existence of *pure-strategy equilibrium* in certain random-player games with independent types are given in the next theorem. Cf. Theorem 4 of Milgrom and Weber (1985).

THEOREM 4. *Suppose there is a partition of S into a finite number of closed sets F_1, F_2, \dots, F_m , such that, for every fixed S , $u(s, S)$ is continuous in s and, for every fixed s , $u(s, S)$ depends only on $S(F_1), S(F_2), \dots, S(F_m)$. If T is a simple mixed sample process, then a pure-strategy equilibrium exists.*

Establishing the existence of correlated equilibrium is conceptually straightforward and does not require any restrictions on \mathcal{T} . If, for every realization of the random-player game \mathcal{T} , the players are assigned strategies according to some correlated (or mixed-strategy) equilibrium of the corresponding *finite*-player game, then players never have any incentive to deviate, since they would have no such incentive even if they knew the number and types of the other players. The following theorem shows that, under a mild continuity condition, a correlated equilibrium can indeed be constructed in such a manner. Essentially, the condition is that the restriction of the payoff function to the strategies of any *finite* number of player types is continuous.

THEOREM 5. *Suppose that, for every finite set B in \mathcal{T} , the restriction of u to the set of all pairs (s, S) such that s belongs to $\tau^{-1}(B)$ and S is supported in $\tau^{-1}(B)$ is continuous. Then, a correlated equilibrium exists.*

It follows from Theorem 3 that every extended Poisson game has a mixed-strategy equilibrium. The same result was proved, with somewhat less generality, by Myerson (1998b, Theorem 1, and 2000, Theorem 0). The existence of mixed-strategy equilibrium in every game with population uncertainty (Myerson, 1998a, Theorem 3) follows from Proposition 3 under the assumption that the number of players has finite expectation. The conditions in Theorem 4 are satisfied by any extended Poisson game in which the player types are selected by the same probability measure λ on \mathcal{T} in all states of the world (which may thus differ only in the expected *number* of players) and λ is nonatomic. Therefore, this theorem shows that such a game has a pure-strategy equilibrium. Theorem 5 implies, in particular, that every random-player game with a continuous payoff function has a correlated equilibrium.

6. ON THE DEFINITION OF CORRELATED EQUILIBRIUM

6.1. Games with a fixed number of players

For n -player *complete* information games, there are several ways correlated equilibrium can be defined. However, these definitions are all equivalent in the sense that they induce the same set of *distributions* of action profiles (Aumann, 1987). One definition employs a set of possible states of the world Ω , with a common prior probability measure on it and n information partitions. The information partitions describe the information available to each player in each state of the world, e.g., the private and public signals he observes. In addition, each state of the world also includes a recommendation of a particular *action* to each player, which the player is assumed to know. The whole structure is said to constitute a correlated equilibrium if, in each state of the world ω , it is Bayes rational for each player to use the recommended action, in the sense that it maximizes his expected payoff with respect to his posterior in ω on Ω . A second, “canonical”, formulation of correlated equilibrium does not involve any information or signals other than the recommended actions. These actions constitute a correlated equilibrium if, whenever an action s is recommended to a player i , it maximizes his expected payoff with respect to the conditional distribution of the other players’ actions, given that i ’s action is s . Formally, this formulation is a special case of the previous one, and corresponds to an information structure in which the states of the world are in one-to-one correspondence with the possible action profiles. Nevertheless, as indicated above, the two formulations are equivalent.

The literature on n -player *incomplete* information games offers several alternative definitions of correlated equilibrium, which are not equivalent to one another. The main differences among them concern the nature of the dependence between the players' types and the signals they observe. In a *strategy correlated equilibrium*, signals and types are independent. In a *type correlated equilibrium*, each player's signal may only depend on his own type. In a *communication equilibrium*, the signals depend on all the players' *reported* types. Finally, in a "*Bayesian solution*" (Forges, 1993), the signals may depend on all the players' true types. The latter concept, which includes the previous ones as special cases, is essentially identical to correlated equilibrium as defined in the previous section. For a complete description of these various notions and their interrelationships, see Forges (1993). The following short account is mainly concerned with how suitable each of the first three notions is as an alternative definition of correlated equilibrium in games with a random number of players. Another alternative definition is discussed in Section 6.2.

In a *communication equilibrium* (e.g., Myerson, 1991), players confidentially report their types to a mediator who, after receiving the reports, confidentially recommends a strategy to each of them. These recommendations may depend on the reported types of all the players in either a deterministic or a random fashion. The equilibrium condition is that the recommendations should be incentive compatible in the sense that it is optimal for all players (i) to report their types honestly and (ii) to follow the mediator's recommendations.

A generalization of communication equilibrium to random-player games seems possible but not straightforward. If the mediator has no sources of information other than the players' reports, then in particular he depends on them to know how many players there are. But if players can choose whether they want their participation to be known to the mediator, the incentive compatibility conditions should be augmented by the requirement that players should never have any incentive to conceal it. This condition differs from (i) above in several ways. For example, a player may conceal his participation from the mediator, but if he chooses not to do so, and to report his presence, the truthfulness of this report is self-evident.

In a *strategy correlated equilibrium* (e.g., Cotter, 1991), a referee, who does not know the players' actual types, confidentially recommends to each of them a pure- or mixed-strategy *profile*, which is an assignment of a pure or mixed strategy to *each* of the player's types. These recommendations are (by construction) independent of the players' actual types, but not necessarily of one another. The equilibrium condition is that it should always be optimal for each player to follow the referee's recommendation by using the pure or mixed strategy corresponding to his *actual* type, assuming that all the other players do the same.

The usefulness of this definition for random-player games is limited. In such a game, not only does the set of players vary in size, but it is also not ordered (i.e., arranged according to players 1, 2, etc.). Therefore, the referee would presumably have to recommend the same strategy profile to all players. In this case, the equilibrium condition reduces to the condition that the recommended strategy profile is always a pure- or mixed-strategy equilibrium, and so the concept of strategy correlated equilibrium reduces to a mixture of such equilibria.

In a *type correlated equilibrium* (e.g., Cotter, 1994), a referee, who does not know the players' actual types, chooses for each of them a pure or mixed strategy for *each* of that player's types. These choices are (by construction) independent of the players' actual types, but not necessarily of one another. In this respect, strategy and type correlated equilibria are similar.

However, in a type correlated equilibrium, each player only gets to know the pure or mixed strategy chosen for *his actual type*. The equilibrium condition is that it should always be optimal for him to use that strategy, assuming that all the other players do the same. This makes strategy correlated equilibrium essentially a special case of type correlated equilibrium.

Myerson (1991, p. 262) finds the assumptions of type correlated equilibrium unnatural. Players are not required to report their types, and yet they learn only what they need to know, namely, their own type's strategy. However, as noted by Cotter (1994), for this to hold, it is not necessary to give different information to different player types; it suffices that they perceive it differently. For example, if two player types, one colorblind and the other nearsighted, are presented with the same visual cue, each is likely to extract different information from it.

6.2. The conditional independence property

In a type (or strategy) correlated equilibrium, if a player's type is known, his strategy does not give any additional information about the other players' types (although it may give information about their strategies). In other words, a player's strategy is conditionally independent of the other players' types, given the player's own type. Proposition 3 of Forges (1993) claims that, in n -player Bayesian games, this *conditional independence property* is characteristic of type correlated equilibria. That is, a correlated equilibrium (in the sense of the definition in Section 5, which is essentially Forges's notion of a Bayesian solution) can be presented as a type correlated equilibrium if and only if the joint distribution of the players' types and their strategies has this property. Cotter (1994) makes a similar claim. However, as the following example shows, these assertions are incorrect.

Example 6. Two players are engaged in a symmetric 2×2 incomplete information game. Each of them either has type t or type t' , with all four type combinations equally likely. The two players always have equal payoffs. If the type combination is (t, t) , their common payoff is given by the following payoff matrix:

$$A = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

For any other type combination, the payoff matrix is $-A$. Consider the following joint distribution of type combinations and strategies. With probability 0.25, the type combination is (t, t) , and in this case, both players always use the same strategy, which has equal probability of being either of the two possible ones. Each of the other three type combinations also has probability 0.25, and in each of them, the two players always use *different* strategies, with the two possible strategy profiles equally probable. Since, with this joint distribution of player types and strategies, both players always receive their highest possible payoff (namely, 1), they do not have any incentive to deviate. However, although the joint distribution is easily seen to satisfy the conditional independence property, it is not generated by any type correlated equilibrium. In fact, in any type correlated equilibrium, both players' expected payoffs must be less than 1. To see that, imagine a referee assigning a mixed strategy to each type of each player in some deterministic or random fashion, *without knowing the players' actual types*. Consider any quartet of mixed strategies $(\mu_{t1}, \mu_{t'1}, \mu_{t2}, \mu_{t'2})$ such that, with positive probability, player 1's types t and t' are assigned the mixed strategies μ_{t1} and $\mu_{t'1}$,

respectively, and player 2's types t and t' the mixed strategies μ_{t2} and $\mu_{t'2}$, respectively. When the type combination is (t, t) , the players' expected payoffs are 1 if and only if there is probability 0 that they use different pure strategies. A necessary condition for this is that there is some pure strategy to which both μ_{t1} and μ_{t2} assign probability 1. For a similar reason, $\mu_{t'1}$ must assign probability 0 to that strategy, and the same must be true for $\mu_{t'2}$. But this implies that both players choose the other, and hence the same, pure strategy when the type combination is (t', t') , and their payoffs in this case are therefore -1 . This proves that the expected payoffs in any type correlated equilibrium are less than 1.

This example shows that the class of correlated equilibria for which the conditional independence property holds is strictly larger than the class of all type (and strategy) correlated equilibria. For random-player games, it is natural to define the former class as consisting of all correlated equilibria in which a player's strategy is conditionally independent of the *number* and types of the other players, given the player's own type. Requiring the conditional independence property gives another possible definition for "correlated equilibrium in incomplete information games".⁵ However, as shown in the next subsection, the class of all correlated equilibria with this property, as well as the classes of strategy correlated equilibria, type correlated equilibria, and communication equilibria, all fail to satisfy a particular condition that "correlated equilibria" may reasonably be required to satisfy, namely, closedness under identification of information types.

6.3. Information types

In n -player *complete* information games, correlated equilibria are equivalent to equilibria with *local interactions* (Mailath *et al.*, 1997), which are defined as follows. For each player i in a finite game G , there is a finite population of agents who could fill i 's role in the game. These agents are matched with those in the populations representing the other players, possibly in a non-uniform way. (For example, some agents may always be matched with the same companions.) An equilibrium with local interactions is a pure-strategy equilibrium in this large game, in which the players are the agents. It is easy to show that any equilibrium with local interactions is transformed into a correlated equilibrium in G when the agents in each of the populations representing the players in G are identified, and their actions are seen as a random choice of action profile for these players. Conversely, every correlated equilibrium in G can be obtained in this manner from some equilibrium with local interaction.

Identification of agents in the populations representing the players in G generally involves loss of information. Specifically, the information an agent's identity gives about the identities of his playmates is lost. Different agents in the same population may be viewed as belonging to different *information types*. The differences between them are payoff-irrelevant, since

⁵ An even stronger requirement than the conditional independence property is that, if a player's type is known, his strategy does not add any information about the other players' *play* (and, in particular, their number and types). In other words, a player's strategy is conditionally independent of the others' play, given the player's type. However, this requirement does not lead to a new kind of correlated equilibrium. Indeed, by Corollary 3, the only correlated strategies with this property are those representing mixed-strategy profiles. Therefore, the only such correlated equilibria are mixed-strategy equilibria.

payoffs are *directly* affected only by what the players do, not by what they know. Identification of information types transforms any correlated equilibrium (and, as a special case, any pure-strategy equilibrium) into a correlated equilibrium in another game. This is because, if a player finds the recommendation to use a particular action acceptable under two different sets of circumstances, he would also find it acceptable if he only knew that one of these sets obtains, but did not know *which* one. Conversely, as indicated above, every correlated equilibrium can be obtained in this way from some pure-strategy equilibrium. This shows that, in games with complete information, the class of correlated equilibria is completely characterized by its closedness under identification of information types and the fact that it includes all pure-strategy equilibria.

In games with *incomplete* information, differences between player types may or may not be payoff-relevant. Yet, it makes sense to require, as in the complete information case, that the set of correlated equilibria be closed under identification of types that do *not* differ from one another in any payoff-relevant way. It is shown below that this requirement essentially leads to the definition of correlated equilibrium given in Section 5.

Identification of information types in random-player games can be formalized as follows. If T and \hat{T} are two random-player games, with spaces of strategies \mathcal{S} and $\hat{\mathcal{S}}$, payoff functions u and \hat{u} , type maps τ and $\hat{\tau}$, and type spaces \mathcal{T} and $\hat{\mathcal{T}}$, then \hat{T} can be obtained from T by *identification of information types* if there is a continuous function $\eta : \mathcal{S} \rightarrow \hat{\mathcal{S}}$ such that:

- (i) $\hat{u}(\hat{s}, \hat{S}) = u(s, S)$ whenever $\hat{s} = \eta(s)$ and $\hat{S} = S \circ \eta^{-1}$;
- (ii) for every $t \in \mathcal{T}$, there is some $\hat{t} \in \hat{\mathcal{T}}$ such that $\eta(\tau^{-1}(\{t\})) = \hat{\tau}^{-1}(\{\hat{t}\})$;⁶ and
- (iii) $\hat{T} = T \circ \theta^{-1}$, where $\theta : \mathcal{T} \rightarrow \hat{\mathcal{T}}$ is defined by $\theta(t) = \hat{\tau}(\eta(t))$.

The function η identifies one or more strategies in the first random-player game with a single strategy in the second game, and hence with one another. Condition (i), which can be written more explicitly as

$$\hat{u}(\eta(s_1), \delta_{\eta(s_2)} + \cdots + \delta_{\eta(s_n)}) = u(s_1, \delta_{s_2} + \cdots + \delta_{s_n}),$$

for every finite list of strategies s_1, s_2, \dots, s_n ($n \geq 1$), says that the strategies thus identified do not differ from one another in any payoff-relevant way. This corresponds, in a complete information game, to the representation of two (strategically) equivalent strategies by a single strategy in the reduced normal form of the game. Condition (ii) refers to strategy sets. It requires the *entire* strategy set of each player type in the first random-player game to be identified with that of a single player type in the second game. In effect, this identifies the first player type, t , with the second one, \hat{t} , and hence also with all the other player types in the first game that are also identified with \hat{t} . Condition (iii) says that the finite point process $T \circ \theta^{-1}$ on $\hat{\mathcal{T}}$ obtained from T through this identification, θ , coincides with \hat{T} . Thus, identifying information types preserves payoffs, maps strategy sets onto strategy sets, and, by means of the induced map between player types, transforms the first random-player game into the

⁶ Obviously, $\hat{t} = \hat{\tau}(\eta(t))$.

second. If the function $\eta : \mathcal{S} \rightarrow \hat{\mathcal{S}}$ is one-to-one and onto and has a continuous inverse, conditions similar to (i)–(iii) are satisfied by the inverse function $\eta^{-1} : \hat{\mathcal{S}} \rightarrow \mathcal{S}$. In this case, the two random-player games may be said to be *isomorphic*.

The next proposition shows that, when one random-player game is transformed into another by means of identification of information types, every correlated equilibrium in the first game is transformed into a correlated equilibrium in the second one. Thus, the class of correlated equilibria, as defined in Section 5, is closed under identification of information types.

PROPOSITION 4. *Let \mathbf{T} be random-player game, $\hat{\mathbf{T}}$ another random-player game, which can be obtained from the first by identification of information types, and $\eta : \mathcal{S} \rightarrow \hat{\mathcal{S}}$ a continuous function satisfying conditions (i)–(iii) above. For every correlated strategy \mathbf{S} in the first random-player game, $\hat{\mathbf{S}} = \mathbf{S} \circ \eta^{-1}$ is a correlated strategy in the second one. If \mathbf{S} is a correlated equilibrium, then so is $\hat{\mathbf{S}}$.*

It follows, in particular, from Proposition 4 that pure-strategy equilibria, which are a special kind of correlated equilibrium, remain correlated equilibria when information types are identified. The next proposition establishes the converse of this: A correlated strategy is a correlated equilibrium *only* if it can be obtained from some pure-strategy equilibrium in another random-player game by identifying information types in a particular way.

PROPOSITION 5. *For every correlated equilibrium $\hat{\mathbf{S}}$ in a random-player game $\hat{\mathbf{T}}$ there is a pure-strategy equilibrium \mathbf{S} in some other random-player game \mathbf{T} , and a continuous function η satisfying conditions (i)–(iii), such that $\hat{\mathbf{S}} = \mathbf{S} \circ \eta^{-1}$.*

Proposition 5 shows that the class of correlated equilibria, as defined in this paper, is the *smallest* one containing the pure-strategy equilibria that is closed under identification of information types. In particular, the classes of type correlated equilibria, and those with the conditional independence property, are not closed under identification of information types. The following example illustrates this.

Example 7. This is a simpler version of Example 6. The type combination is either (t, t) or (t', t') , but not any of the other two. The players' payoffs are 0 regardless of their types and actions. Consider the pure-strategy equilibrium in which each player uses the first or second strategy when his type is t or t' , respectively. This is clearly a correlated equilibrium according to all the above definitions. However, if player 1's two types (which, of course, differ from one another only in payoff-irrelevant ways) are identified, then this correlated equilibrium is transformed into a correlated strategy that is not a type or strategy correlated equilibrium. Indeed, the conditional independence property does not hold, since player 1's action *does* reveal player 2's type.

7. SUBJECTIVE BELIEFS

The definition of correlated and other kinds of equilibria in this paper is based on a notion of *objective* uncertainty, which stems from the random nature of the game. Players act as rational Bayesian decision makers, whose beliefs about the number and types of the other players are conditional probabilities, derived from a common prior on the number and types of players. It

is possible to extend this model by allowing the players to hold arbitrary *subjective* beliefs about the game, the other players' beliefs about it, and so on. In this extension, which is outlined in this section, a player's type is a composite entity, comprising two components: the first specifying the player's physical, payoff-relevant attributes; and the second his first-order beliefs about the number and (composite) types of the other players as a finite point process on the space of all player types. From these first-order beliefs, and those of the other player types, it is possible to construct the complete hierarchy of beliefs of each type. This formulation closely parallels the models of Mertens and Zamir (1985) and Brandenburger and Dekel (1993), which, however, are only applicable to situations involving a commonly known number of players. Here, by contrast, it only has to be commonly known that the set of players is finite.

The space of all "basic types" of players is a Polish space \mathcal{T} . The interpretation of a basic type is that it specifies all payoff-relevant information about the player, so that a finite point measure on \mathcal{T} completely specifies the game. A player's first-order beliefs about the number of other players and their basic types are described by a finite point process on $\mathcal{X}_0 = \mathcal{T}$. The basic type of a player and his first-order beliefs about the number and basic types of the other players are, therefore, jointly described by an element of the product space $\mathcal{X}_1 = \mathcal{T} \times \Delta(\mathcal{N}_{\mathcal{T}})$. Proceeding inductively, the (Polish) space \mathcal{X}_i of the players' basic types and their i -order beliefs ($i = 1, 2, \dots$) is defined by $\mathcal{X}_i = \mathcal{T} \times \Delta(\mathcal{N}_{\mathcal{X}_{i-1}})$.

A player's high-order beliefs should, of course, be consistent with his lower-order ones. To express this *coherency* requirement formally, define $\xi_0 : \mathcal{X}_1 \rightarrow \mathcal{X}_0$ as the projection on the first coordinate, and proceed inductively to define a continuous function ξ_i from \mathcal{X}_{i+1} onto⁷ \mathcal{X}_i ($i = 1, 2, \dots$) as follows: If t is an element of \mathcal{T} and \mathbf{X}_i is a finite point process on \mathcal{X}_i , then $\xi_i(t, \mathbf{X}_i) = (t, \mathbf{X}_i \circ \xi_{i-1}^{-1})$. A *player type* is an element $x = (x_0, x_1, \dots)$ of $\mathcal{X}_0 \times \mathcal{X}_1 \times \dots$ such that, for all $i \geq 0$, $x_i = \xi_i(x_{i+1})$. The space of all player types is denoted by \mathcal{X} . Since it is a closed subset of the product space $\mathcal{X}_0 \times \mathcal{X}_1 \times \dots$, this space is Polish in the relative (product) topology. As a measurable space, \mathcal{X} is the *inverse limit* of $\mathcal{X}_0, \mathcal{X}_1, \dots$ relative to the maps ξ_0, ξ_1, \dots (Parthasarathy, 1967, Theorem V.3.2).

There is a natural one-to-one correspondence—indeed a homeomorphism—between the space of all player types \mathcal{X} and the set $\mathcal{T} \times \Delta(\mathcal{N}_{\mathcal{X}})$ of all pairs consisting of a basic type and a finite point process on \mathcal{X} . Thus, specifying a player's basic type and his hierarchy of beliefs is equivalent to specifying the player's basic type and his beliefs about the number, basic types, and hierarchies of beliefs of the *other* players. This homeomorphism, $f : \mathcal{X} \rightarrow \mathcal{T} \times \Delta(\mathcal{N}_{\mathcal{X}})$, is constructed below.

For all $i \geq 0$, \mathcal{X}_i is a Polish space and ξ_i is continuous and onto. Therefore, by Theorem V.3.2 of Parthasarathy (1967), for every sequence X_0, X_1, \dots such that, for all $i \geq 0$, X_i is a finite

⁷ The fact that ξ_i is *onto* \mathcal{X}_i follows from the existence, for every $i \geq 0$, of a continuous function $\theta_{i+1} : \mathcal{X}_i \rightarrow \mathcal{X}_{i+1}$ such that $\xi_i \circ \theta_{i+1}$ is the identity map on \mathcal{X}_i . The first of these functions is defined by $\theta_1(t) = (t, \mathbf{T}^*)$, where \mathbf{T}^* is some fixed finite point process on \mathcal{T} . Proceeding inductively, for all $i \geq 1$, $\theta_{i+1}(t, \mathbf{X}_{i-1}) = (t, \mathbf{X}_{i-1} \circ \theta_i^{-1})$.

point measure on \mathcal{X}_i and $X_i = X_{i+1} \circ \xi_i^{-1}$, there is a unique measure X on (the inverse limit) \mathcal{X} such that $X_i = X \circ \pi_i^{-1}$ for all i , where $\pi_i : \mathcal{X} \rightarrow \mathcal{X}_i$ is the projection on the i th coordinate. It is, moreover, not very difficult to see that: (i) X is a finite point measure; (ii) the function $(X_0, X_1, \dots) \mapsto X$ from the subspace of $\mathcal{N}_{\mathcal{X}_0} \times \mathcal{N}_{\mathcal{X}_1} \times \dots$ consisting of all sequences as above to $\mathcal{N}_{\mathcal{X}}$ is one-to-one and onto; and (iii) both this function and its inverse are continuous. It follows, again by the same theorem, that if X_0, X_1, \dots are such that, for all $i \geq 0$, X_i is a finite point process on \mathcal{X}_i and $X_i = X_{i+1} \circ \xi_i^{-1}$, then there is a finite point process \mathbf{X} on \mathcal{X} such that $X_i = \mathbf{X} \circ \pi_i^{-1}$ for all i (and if \mathbf{X}' is another such finite point process, then it is equal in distribution to \mathbf{X}). Therefore, for every element $x = (x_0, x_1, \dots)$ of \mathcal{X} , there is a unique finite point process \mathbf{X} on \mathcal{X} such that $x_i = (x_0, \mathbf{X} \circ \pi_{i-1}^{-1})$ for all $i \geq 1$. Since the cylinders in $\mathcal{N}_{\mathcal{X}_0} \times \mathcal{N}_{\mathcal{X}_1} \times \dots$ constitute a convergence-determining class (see Billingsley, 1999, Theorem 2.4), the function $x \mapsto (x_0, \mathbf{X})$ is continuous. This function from \mathcal{X} to $\mathcal{T} \times \Delta(\mathcal{N}_{\mathcal{X}})$, which will be denoted by f , has a continuous inverse, $g : \mathcal{T} \times \Delta(\mathcal{N}_{\mathcal{X}}) \rightarrow \mathcal{X}$, which is defined as follows: If t is an element of \mathcal{T} and \mathbf{X} is a finite point process on \mathcal{X} , then $g(t, \mathbf{X}) = (t, x_1, x_2, \dots)$, where $x_i = (t, \mathbf{X} \circ \pi_{i-1}^{-1})$ for all $i \geq 1$. Therefore, f maps \mathcal{X} *homeomorphically* onto $\mathcal{T} \times \Delta(\mathcal{N}_{\mathcal{X}})$.

In the model described in this section, for every basic type and (coherent) hierarchy of beliefs, there is a corresponding player type with the same basic type and beliefs about the other players. Herein lies the basic difference between this model and the one described in Section 3, in which, for every player type (a notion corresponding to “basic type” here) there is just one hierarchy of beliefs, namely, that induced by the random-player game \mathbf{T} . These may be called the *objective* beliefs determined by \mathbf{T} . The objective beliefs define a one-to-one measurable function $\rho : \mathcal{T} \rightarrow \mathcal{X}$. Specifically, if ρ_0 is defined as the identity map on \mathcal{T} and $t \mapsto \mathbf{T}_t$ is the map from \mathcal{T} to $\Delta(\mathcal{N}_{\mathcal{T}})$ defined by Eq. (1), then $\rho(t) = (\rho_0(t), \rho_1(t), \dots)$ for all $t \in \mathcal{T}$, where, for all $i \geq 1$, the measurable function $\rho_i : \mathcal{T} \rightarrow \mathcal{X}_i$ is defined inductively by $\rho_i(t) = (t, \mathbf{T}_t \circ \rho_{i-1}^{-1})$.⁸ (That $\rho(t)$ indeed lies in \mathcal{X} follows from the equality $\xi_i \circ \rho_{i+1} = \rho_i$, which can easily be proved by induction for all $i \geq 0$.) Thus, fixing the random-player game and insisting that players have objective beliefs effectively excludes all player types not belonging to the (measurable) set $\rho(\mathcal{T})$. The random-player game \mathbf{T} may be identified with the finite point process $\mathbf{X} = \mathbf{T} \circ \rho^{-1}$ on \mathcal{X} .⁹ It is not difficult to prove that, for $\mathbb{E}\mathbf{X}$ -almost all player types x , the beliefs of that player type about the number and types of the other players coincide with his posterior \mathbf{X}_x as defined by Eq. (1), with \mathbf{X} and x replacing \mathbf{T} and t , respectively. Formally, $f(x) = (x_0, \mathbf{X}_x)$. The finite point process \mathbf{X} is thus *consistent* (cf. Mertens and Zamir, 1985, Section 4) in the sense that the beliefs of almost all players are (a version of) the conditional distribution of this finite point process.

⁸ An equivalent implicit definition for ρ is $f(\rho(t)) = (t, \mathbf{T}_t \circ \rho^{-1})$, for all $t \in \mathcal{T}$. (The function $f : \mathcal{X} \rightarrow \mathcal{X} \times \Delta(\mathcal{N}_{\mathcal{X}})$ is the homeomorphism defined in the previous paragraph.)

⁹ Note that \mathbf{X} only selects elements of $\rho(\mathcal{T})$. And, by the formula in footnote 8, a player whose type belongs to $\rho(\mathcal{T})$ believes that the types of all the other players are also in that set. Thus, in a sense, it is common knowledge among the players that their types are all in $\rho(\mathcal{T})$.

APPENDIX

The appendix presents the proofs of the various propositions and theorems in this paper.

Proof of Theorem 1. For a proof of this theorem, see Kallenberg (1986, Theorem 11.5 and Exercise 11.1). Alternatively, the theorem can easily be deduced from Theorem 2 and Propositions 1 and 2. ■

Proof of Proposition 1. If the distribution of the random variable $T_t(\mathcal{J})$ is the same for $\mathbb{E}T$ -almost all t , then it follows from Eq. (1) that, for all $n \geq 1$ and all measurable sets B ,

$$(10) \quad \mathbb{P}(T_t(\mathcal{J}) = n - 1) \mathbb{E}T(B) = \mathbb{E}[1_{T(\mathcal{J})=n} T(B)].$$

For $B = \mathcal{J}$, this gives (i). If B is such that, for some integer n_0 , $T(B) \leq n_0$ almost surely, then (10) implies $\mathbb{P}(T_t(\mathcal{J}) = n - 1) \mathbb{E}T(B) \leq n_0 \mathbb{P}(T(\mathcal{J}) = n)$. It then follows by comparison with (3) that, for every n such that $\mathbb{P}(T(\mathcal{J}) = n) > 0$, $n \mathbb{E}T(B) \leq n_0 \mathbb{E}T(\mathcal{J})$. Therefore, $T(\mathcal{J})$ is bounded (i.e., not unbounded) or $\mathbb{E}T(B) = 0$. This proves (ii). The special case in which B is a singleton gives (iii). If $T_t(\mathcal{J}) = T(\mathcal{J})$, then it follows from (3) that $T(\mathcal{J})$ has a Poisson distribution, and is therefore unbounded unless $\mathbb{E}T(\mathcal{J}) = 0$. In view of (iii), this proves the last part of the proposition. ■

Proof of Theorem 2 (an outline). Fix some $n \geq 1$ such that $\mathbb{P}(S(\mathcal{S}) = n) > 0$. There exists a symmetric probability measure on \mathcal{S}^n , and n random elements s_1, s_2, \dots, s_n of \mathcal{S} whose joint distribution is given by that measure, such that the conditional distribution of S , given that $S(\mathcal{S}) = n$, is equal to the distribution of $\delta_{s_1} + \delta_{s_2} + \dots + \delta_{s_n}$. For $\mathbb{E}S$ -almost all s such that $\mathbb{P}(S_s(\mathcal{S}) = n - 1) > 0$, the conditional distribution of S_s , given that $S_s(\mathcal{S}) = n - 1$, is equal to the conditional distribution of $\delta_{s_2} + \dots + \delta_{s_n}$, given that $s_1 = s$. Suppose that, for every $t \in \mathcal{T}$, there is a finite point process S_t on \mathcal{S} , such that (6) holds for every s . Then, s_2, \dots, s_n are conditionally independent of s_1 , given $\tau(s_1)$, and hence also given $\tau(s_1), \tau(s_2), \dots, \tau(s_n)$. Since there is obviously no special role for the first random element, s_1, s_2, \dots, s_n are conditionally independent, given $\tau(s_1), \tau(s_2), \dots, \tau(s_n)$. This implies that the distribution of $\delta_{s_1} + \delta_{s_2} + \dots + \delta_{s_n}$ is equal to that of the finite point process obtained from $\delta_{\tau(s_1)} + \delta_{\tau(s_2)} + \dots + \delta_{\tau(s_n)}$ through the transition probability $t \mapsto \mu_t^{(n)}$, where $\mu_t^{(n)}$ denotes the conditional distribution of s_1 , given that $\tau(s_1) = t$. Therefore, to prove that S can be obtained from the finite point process $S \circ \tau^{-1}$ on \mathcal{T} through some transition probability, it suffices to show that $\mu_t^{(n)}$ does not, in fact, depend on n .

Denote $S \circ \tau^{-1}$ by T . Since $\mathbb{E}S \circ \tau^{-1} = \mathbb{E}T$, and $\tau: \mathcal{S} \rightarrow \mathcal{T}$ is a measurable function from one Polish space onto another, there exists a transition probability $t \mapsto \mu_t$ such that (5) and (7) hold (Parthasarathy, 1967, Theorem V.8.1). By these equations and (6), for every $f \in \mathfrak{F}(\mathcal{S})$, on the one hand,

$$\int \mathbb{P}(S_s(\mathcal{S}) = n - 1) f(s) \mathbb{E}S(ds) = \int \mathbb{P}(S_{\tau(s)}(\mathcal{S}) = n - 1) f(s) \mathbb{E}S(ds)$$

$$= \int \mathbb{P}(\mathbf{S}_t(\mathcal{S}) = n-1) \left[\int f(s) \mu_t(ds) \right] \mathbb{E}T(dt),$$

and on the other hand, by Eq. (2),

$$\int \mathbb{P}(\mathbf{S}_s(\mathcal{S}) = n-1) f(s) \mathbb{E}S(ds) = \mathbb{E} \left[\mathbf{1}_{\mathbf{S}(\mathcal{S})=n} \int f(s) \mathbf{S}(ds) \right] = n \mathbb{P}(\mathbf{S}(\mathcal{S}) = n) \mathbb{E}[f(s_1)].$$

Hence,

$$(11) \quad \mathbb{E}[f(s_1)] = \int \left[\int f(s) \mu_t(ds) \right] h(t) \mathbb{E}T(dt) \quad (f \in \mathfrak{F}(\mathcal{S})),$$

where $h(t)$ denotes the quotient $\mathbb{P}(\mathbf{S}_t(\mathcal{S}) = n-1)/[n \mathbb{P}(\mathbf{S}(\mathcal{S}) = n)]$. For every $g \in \mathfrak{F}(\mathcal{T})$, substituting $g(\tau(s))$ for $f(s)$ in (11) gives $\mathbb{E}[g(\tau(s_1))] = \int g(t) h(t) \mathbb{E}T(dt)$. In particular, for $g(t) = \int f(s) \mu_t(ds)$,

$$\mathbb{E} \left[\int f(s) \mu_{\tau(s_1)}(ds) \right] = \int \left[\int f(s) \mu_t(ds) \right] h(t) \mathbb{E}T(dt).$$

Therefore, by (11), $\mathbb{E}[f(s_1)] = \mathbb{E}[\int f(s) \mu_{\tau(s_1)}(ds)]$. This shows that, for every t , μ_t is a version of the conditional distribution of s_1 , given that $\tau(s_1) = t$. Hence, this conditional distribution does not depend on n . As shown above, this proves that the distribution of \mathbf{S} is equal to that of the finite point process obtained from \mathbf{T} through the transition probability $t \mapsto \mu_t$. \blacksquare

Proof of Proposition 2. It follows from (4), the definition of \mathbf{S}_T , and (7) that, for every $g \in \mathfrak{F}(\mathcal{N}_T)$, $\mathbb{E}[g(\mathbf{S} \circ \tau^{-1})] = \mathbb{E}_T[\mathbb{E}_{\mathbf{S}_T}[g(\mathbf{S}_T \circ \tau^{-1})]] = \mathbb{E}[g(\mathbf{T})]$. Therefore, $\mathbf{T} = \mathbf{S} \circ \tau^{-1}$, which shows that \mathbf{T} can be obtained from \mathbf{S} through some transition probability, namely, $s \mapsto \delta_{\tau(s)}$. Since, by assumption, the converse holds as well, \mathbf{S} is a mixed sample process if and only if \mathbf{T} is a mixed sample process, and in this case, the distribution of the sample size is the same for both. It remains to prove (i).

It follows from (4), the definition of \mathbf{S}_T , (2), (4) again (with \mathbf{S}_t and \mathbf{T}_t replacing \mathbf{S} and \mathbf{T} , respectively), (7), and (5) that, for every $h \in \mathfrak{F}(\mathcal{S} \times \mathcal{N}_S)$,

$$\begin{aligned} \mathbb{E} \left[\int h(s, \mathbf{S} - \delta_s) \mathbf{S}(ds) \right] &= \mathbb{E}_T \left[\mathbb{E}_{\mathbf{S}_T} \left[\int h(s, \mathbf{S}_T - \delta_s) \mathbf{S}_T(ds) \right] \right] \\ &= \mathbb{E}_T \left[\int \mathbb{E}_{\mathbf{S}_{T-\delta_t}} \left[\int h(s, \mathbf{S}_{T-\delta_t}) \mu_t(ds) \right] \mathbf{T}(dt) \right] = \int \mathbb{E}_{T_t} \left[\mathbb{E}_{\mathbf{S}_{T_t}} \left[\int h(s, \mathbf{S}_{T_t}) \mu_t(ds) \right] \right] \mathbb{E}T(dt) \\ &= \int \mathbb{E} \left[\int h(s, \mathbf{S}_t) \mu_t(ds) \right] \mathbb{E}T(dt) = \int \int \mathbb{E} \left[h(s, \mathbf{S}_{\tau(s)}) \right] \mu_t(ds) \mathbb{E}T(dt) = \int \mathbb{E} \left[h(s, \mathbf{S}_{\tau(s)}) \right] \mathbb{E}S(ds). \end{aligned}$$

Therefore, comparison with (2) shows that (6) holds for $\mathbb{E}S$ -almost all s . \blacksquare

The proof of Theorem 3 requires four lemmas. The first lemma gives an ‘‘ex ante’’ characterization of correlated equilibrium in a random-player game \mathbf{T} in terms of the expected aggregate utility (which is the expression on the left-hand side of Eq. (12) below).

LEMMA 1. *Suppose that, for every play S and every player type t , the restriction of $u(\cdot, S)$ to the strategy set of t is continuous. Then, a correlated strategy \mathbf{S} is a correlated equilibrium if and only if*

$$(12) \quad \mathbb{E}\left[\int u(s, \mathbf{S} - \delta_s) \mathbf{S}(ds)\right] = \sup_{\phi} \mathbb{E}\left[\int u(\phi(s), \mathbf{S} - \delta_s) \mathbf{S}(ds)\right],$$

where the supremum is taken over the set of all measurable functions $\phi : \mathcal{S} \rightarrow \mathcal{S}$ such that $\tau \circ \phi = \tau$.

Proof. The set $\{(s', s) \mid \tau(s') = \tau(s)\}$ is closed in $\mathcal{S} \times \mathcal{S}$, and its sections are compact and nonempty. The restriction of the function $(s', s) \mapsto \mathbb{E}[u(s', \mathbf{S}_s)]$ to this set is measurable, and is continuous in s' for every fixed $s \in \mathcal{S}$. Since \mathcal{S} is Polish, these facts guarantee the existence of a measurable function $\phi_0 : \mathcal{S} \rightarrow \mathcal{S}$ such that, for every s , $\tau(\phi_0(s)) = \tau(s)$ and $\mathbb{E}[u(\phi_0(s), \mathbf{S}_s)] = \max_{s' \in \tau^{-1}(\{\tau(s)\})} \mathbb{E}[u(s', \mathbf{S}_s)]$ (Wagner, 1977, Theorem 9.2). Since, by Eq. (2), $\mathbb{E}\left[\int u(\phi(s), \mathbf{S} - \delta_s) \mathbf{S}(ds)\right] = \int \mathbb{E}[u(\phi(s), \mathbf{S}_s)] \mathbb{E}\mathbf{S}(ds)$ for every measurable function $\phi : \mathcal{S} \rightarrow \mathcal{S}$, the supremum on the right-hand side of (12) is attained at ϕ_0 . Therefore, (12) is equivalent to the equality

$$\int \mathbb{E}[u(s, \mathbf{S}_s)] \mathbb{E}\mathbf{S}(ds) = \int \max_{s' \in \tau^{-1}(\{\tau(s)\})} \mathbb{E}[u(s', \mathbf{S}_s)] \mathbb{E}\mathbf{S}(ds).$$

This equality clearly holds if and only if \mathbf{S} is a correlated equilibrium. ■

LEMMA 2. *The distribution of every extended sample process is absolutely continuous with respect to the distribution of some finite Poisson process.*

Proof. Let \mathcal{S} and \mathcal{T} be Polish spaces, and \mathbf{S} a finite point process on \mathcal{S} that is obtained from a finite point process \mathbf{T} on \mathcal{T} through a transition probability $t \mapsto \mu_t$ with countable range. Let $\mathcal{T}' = \{t_1, t_2, \dots\}$ be a finite or countably infinite subset of \mathcal{T} such that the range of the above transition probability is included in $\{\mu_{t_1}, \mu_{t_2}, \dots\}$. There exists a measurable function $\varphi : \mathcal{T} \rightarrow \mathcal{T}'$ with $\varphi(\mathcal{T}) \subseteq \mathcal{T}'$ such that $\mu_{\varphi(t)} = \mu_t$ for all t . The finite point process \mathbf{S}' on \mathcal{S} obtained from $\mathbf{T} \circ \varphi^{-1}$ through the transition probability $t \mapsto \mu_t$ is clearly equal in distribution to \mathbf{S} . For each $t_i \in \mathcal{T}'$, choose a probability $p_i > 0$, such that $p_1 + p_2 + \dots = 1$. Let \mathbf{P} be a finite Poisson process on \mathcal{T} obtained by independently drawing n elements of \mathcal{T}' according to these probabilities, with n itself determined as a Poisson random variable with some nonzero parameter. For every $T \in \mathcal{N}_{\mathcal{T}}$, $\mathbb{P}(\mathbf{P} = T) > 0$ if and only if T is supported in \mathcal{T}' . The set of all finite point measures on \mathcal{T} satisfying this inequality is therefore countable, and includes all $T \in \mathcal{N}_{\mathcal{T}}$ such that $\mathbb{P}(\mathbf{T} \circ \varphi^{-1} = T) > 0$. Let \mathbf{S}'' be the finite Poisson process on \mathcal{S} that is obtained from \mathbf{P} through the transition probability $t \mapsto \mu_t$. For every $g \in \mathfrak{F}(\mathcal{N}_{\mathcal{S}})$,

$$\mathbb{E}[g(\mathbf{S}'')] = \sum_{T \in \mathcal{N}_{\mathcal{T}}} \mathbb{P}(\mathbf{P} = T) \mathbb{E}[g(\mathbf{S}_T)] \quad \text{and} \quad \mathbb{E}[g(\mathbf{S}')] = \sum_{T \in \mathcal{N}_{\mathcal{T}}} \mathbb{P}(\mathbf{T} \circ \varphi^{-1} = T) \mathbb{E}[g(\mathbf{S}_T)],$$

where \mathbf{S}_T is defined as in Section 2.3. Therefore, $\mathbb{E}[g(\mathbf{S}'')] = 0$ implies that $\mathbb{E}[g(\mathbf{S}')] = 0$. This shows that the distribution of \mathbf{S}' is absolutely continuous with respect to that of \mathbf{S}'' . ■

LEMMA 3. *Suppose that the distribution of a finite point process \mathbf{T} on a Polish space \mathcal{T} is absolutely continuous with respect to the distribution of some extended sample process \mathbf{T}' . Then it is also absolutely continuous with respect to the distribution of every finite Poisson process \mathbf{P} that satisfies $\mathbb{E}\mathbf{P} = \mathbb{E}\mathbf{T}$.*

Proof. In view of Lemma 2, it may be assumed that \mathbf{T}' is a finite Poisson process. By the Radon-Nikodym theorem, there exists a function $g \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})$ such that $\mathbb{E}[h(\mathbf{T})] = \mathbb{E}[h(\mathbf{T}') g(\mathbf{T}')] for all $h \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})$. In particular, for every measurable set B in \mathcal{T} ,$

$$\mathbb{E}\mathbf{T}(B) = \mathbb{E}[\mathbf{T}'(B) g(\mathbf{T}')] = \int_B \mathbb{E}[g(\mathbf{T}'_t + \delta_t)] \mathbb{E}\mathbf{T}'(dt),$$

where the second equality follows from (2). Therefore, $\mathbb{E}\mathbf{T}$ is absolutely continuous with respect to $\mathbb{E}\mathbf{T}'$, and the Radon-Nikodym derivative of $\mathbb{E}\mathbf{T}$ with respect to $\mathbb{E}\mathbf{T}'$ is the function $f \in \mathfrak{F}(\mathcal{T})$ defined by $f(t) = \mathbb{E}[g(\mathbf{T}'_t + \delta_t)]$. Note that $\mathbb{E}[\mathbf{T}'(f^{-1}(\{0\})) g(\mathbf{T}')] = \int_{f^{-1}(\{0\})} f(t) \mathbb{E}\mathbf{T}'(dt) = 0$, and, therefore, it may be assumed that $g(\mathbf{T}) = 0$ for every nonzero finite point measure $\mathbf{T} = \delta_{t_1} + \delta_{t_2} + \dots + \delta_{t_n}$ such that $f(t_i) = 0$ for some $1 \leq i \leq n$. Since \mathbf{T}' is a finite Poisson process, for every $h \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})$,

$$\begin{aligned} \mathbb{E}[h(\mathbf{T})] &= \mathbb{E}[h(\mathbf{T}') g(\mathbf{T}')] = \sum_{n=0}^{\infty} \frac{1}{n!} e^{-\mathbb{E}\mathbf{T}'(\mathcal{T})} \int \dots \int h\left(\sum_{i=1}^n \delta_{t_i}\right) g\left(\sum_{i=1}^n \delta_{t_i}\right) \prod_{i=1}^n \mathbb{E}\mathbf{T}'(dt_i) \\ &= \sum_{n=0}^{\infty} \frac{1}{n!} e^{-\mathbb{E}\mathbf{T}'(\mathcal{T})} \int \dots \int h\left(\sum_{i=1}^n \delta_{t_i}\right) g\left(\sum_{i=1}^n \delta_{t_i}\right) \frac{1}{\prod_{i=1}^n f(t_i)} \prod_{i=1}^n \mathbb{E}\mathbf{T}(dt_i) \\ &= \sum_{n=0}^{\infty} \frac{1}{n!} e^{-\mathbb{E}\mathbf{T}(\mathcal{T})} \int \dots \int h\left(\sum_{i=1}^n \delta_{t_i}\right) \hat{g}\left(\sum_{i=1}^n \delta_{t_i}\right) \prod_{i=1}^n \mathbb{E}\mathbf{T}(dt_i), \end{aligned}$$

where $\hat{g} \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})$ is defined by $\hat{g}(\mathbf{T}) = g(\mathbf{T}) e^{\mathbb{E}\mathbf{T}(\mathcal{T}) - \mathbb{E}\mathbf{T}'(\mathcal{T}) - \int \ln f(t) \mathbf{T}(dt)}$. Therefore, if \mathbf{P} is any finite Poisson process such that $\mathbb{E}\mathbf{P} = \mathbb{E}\mathbf{T}$, then $\mathbb{E}[h(\mathbf{T})] = \mathbb{E}[h(\mathbf{P}) \hat{g}(\mathbf{P})]$, which shows that the distribution of \mathbf{T} is absolutely continuous with respect to that of \mathbf{P} . ■

LEMMA 4. *A measure on the space of strategies is a distributional strategy if and only if it is the mean measure of some correlated strategy. The set of all distributional strategies is compact, and the same is true for the set of all distributions of correlated strategies.*

Proof. The fact that the mean measure of a correlated strategy is a distributional strategy follows immediately from their definitions. The converse is shown by (9). It follows that the continuous function $\mathcal{S} \mapsto \mathbb{E}\mathcal{S}$, which sends each correlated strategy \mathcal{S} to its mean measure, is *onto* the set of distributional strategies. Therefore, to prove that the last-mentioned set is compact, it suffices to show that the set of all distributions of correlated strategies is compact in $\Delta(\mathcal{N}_{\mathcal{S}})$. (A direct proof is also easy to give.) Since this set is easily seen to be closed, it suffices, by Prohorov's theorem (Billingsley, 1999, Theorem 5.1), to show that it is tight. This can be shown as follows. Fix $\varepsilon > 0$. Since \mathcal{T} is a Polish space and the mean measure of the random-player game \mathbf{T} is finite, there is a compact set $K \subseteq \mathcal{T}$ such that $\mathbb{E}\mathbf{T}(\mathcal{T} \setminus K) < \varepsilon$ (Billingsley, 1999, Theorem 1.3). Hence, the compact set $K' = \tau^{-1}(K)$ has the property that $\mu(\mathcal{S} \setminus K') < \varepsilon$ for all distributional strategies μ . Let $K'' \subseteq \mathcal{N}_{\mathcal{S}}$ be the set of all finite point measures \mathcal{S} on \mathcal{S} that satisfy $\mathcal{S}(K') < 1/\varepsilon$ and $\mathcal{S}(\mathcal{S} \setminus K') = 0$. This set is compact (see Section 2). For every correlated strategy \mathcal{S} , $\mathbb{P}(\mathcal{S} \notin K'') \leq \varepsilon \mathbb{E}\mathcal{S}(K') + \mathbb{E}\mathcal{S}(\mathcal{S} \setminus K') < \varepsilon \mathbb{E}\mathbf{T}(\mathcal{T}) + \varepsilon$, since $\mathbb{E}\mathcal{S}$

is a distributional strategy. Since ε can be chosen arbitrarily small, this proves that the set of all distributions of correlated strategies is tight. ■

Proof of Theorem 3. Since the mean measure of \mathbf{T} is finite, there is a finite Poisson process \mathbf{P} on \mathcal{T} such that $\mathbb{E}\mathbf{P} = \mathbb{E}\mathbf{T}$. If the distribution of \mathbf{T} is absolutely continuous with respect to the distribution of some extended sample process then, by Lemma 3 and the Radon-Nikodym theorem, there is some $g \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})$ such that $\mathbb{E}[h(\mathbf{T})] = \mathbb{E}[h(\mathbf{P}) g(\mathbf{P})]$ for all $h \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})$.

Consider a correlated strategy \mathbf{S} obtained from \mathbf{T} through some mixed-strategy profile $t \mapsto \mu_t$, and let $\phi: \mathcal{S} \rightarrow \mathcal{S}$ be some measurable function. By Eq. (4),

$$(13) \quad \mathbb{E}\left[\int u(\phi(s), \mathbf{S} - \delta_s) \mathbf{S}(ds)\right] = \mathbb{E}_{\mathbf{T}}\left[\mathbb{E}_{\mathbf{S}_{\mathbf{T}}}\left[\int u(\phi(s), \mathbf{S}_{\mathbf{T}} - \delta_s) \mathbf{S}_{\mathbf{T}}(ds)\right]\right].$$

By the definitions of $\mathbf{S}_{\mathbf{T}}$ (in Section 2.3) and g , the right-hand side of (13) is equal to

$$\sum_{n=1}^{\infty} q_n \int \int \cdots \int \left[\int \int \cdots \int u(\phi(s_1), \sum_{i=2}^n \delta_{s_i}) \prod_{i=1}^n \mu_{t_i}(ds_i) \right] g\left(\sum_{i=1}^n \delta_{t_i}\right) \prod_{i=1}^n \mathbb{E}\mathbf{P}(dt_i),$$

where $q_n = (1/(n-1!)) e^{-\mathbb{E}\mathbf{P}(\mathcal{T})}$. The assumption $\mathbb{E}\mathbf{P} = \mathbb{E}\mathbf{T}$ implies that this expression is equal to

$$(14) \quad \sum_{n=1}^{\infty} q_n \int \int \cdots \int u(\phi(s_1), \sum_{i=2}^n \delta_{s_i}) g\left(\sum_{i=1}^n \delta_{\tau(s_i)}\right) \prod_{i=1}^n \mu(ds_i),$$

where μ is the distributional strategy defined in Eq. (8). For any two finite measures on \mathcal{S} , or linear combinations of such measures, μ' and μ'' , define

$$G(\mu', \mu'') = \sum_{n=1}^{\infty} q_n \int \int \cdots \int u(s_1, \sum_{i=2}^n \delta_{s_i}) g\left(\sum_{i=1}^n \delta_{\tau(s_i)}\right) \mu'(ds_1) \prod_{i=2}^n \mu''(ds_i).$$

If $\tau = \tau \circ \phi$, then the expression in (14), and hence also that on the left-hand side of (13), are equal to $G(\mu \circ \phi^{-1}, \mu)$. Moreover, $(\mu \circ \phi^{-1}) \circ \tau^{-1} = \mu \circ \tau^{-1} = \mathbb{E}\mathbf{T}$, and, hence, $\mu \circ \phi^{-1}$, like μ , is a distributional strategy. It follows, by Lemma 1, that a sufficient condition for $t \mapsto \mu_t$ to be a mixed-strategy equilibrium in the random-player game \mathbf{T} is that $G(\mu, \mu) \geq G(\mu', \mu)$ for all distributional strategies μ' . In other words, it suffices that (μ, μ) is a symmetric equilibrium in the symmetric two-person game in which each player chooses a distributional strategy and, when one player chooses μ' and the other μ'' , the first player's payoff is $G(\mu', \mu'')$. Therefore, to complete the proof of the theorem it suffices to show that this symmetric two-person game has a symmetric equilibrium.

By Lemma 4, the set \mathcal{M} of all distributional strategies is compact in the weak topology. Since \mathcal{M} can be seen as a convex subset of a locally convex Hausdorff linear topological space (namely, the space of all linear combinations of finite measures on \mathcal{S}), and G is linear in its first argument, it follows from a well-known fixed-point theorem (Fan, 1952) that a sufficient condition for the existence of a symmetric equilibrium is that the restriction of G to $\mathcal{M} \times \mathcal{M}$ is continuous.

The space $\mathcal{N}_{\mathcal{T}}$ is Polish, and the function $g : \mathcal{N}_{\mathcal{T}} \rightarrow \mathbb{R}_+$ satisfies $\mathbb{E}[\mathbf{P}(\mathcal{T}) g(\mathbf{P})] = \mathbb{E}[\mathbf{T}(\mathcal{T})] < \infty$. Therefore, for every $k \geq 1$, there is a bounded continuous function $g_k : \mathcal{N}_{\mathcal{T}} \rightarrow \mathbb{R}_+$ such that $\mathbb{E}[\mathbf{P}(\mathcal{T}) |g(\mathbf{P}) - g_k(\mathbf{P})|] < 1/k$ (see Dunford and Schwartz, 1958, Lemma IV.8.19). Since the function sending two (or more) probability measures on \mathcal{S} to their product is continuous (Billingsley, 1999, Theorem 2.8), the function $G_k : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ defined by

$$G_k(\mu', \mu'') = \sum_{n=1}^{\infty} q_n \int \int \cdots \int u(s_1, \sum_{i=2}^n \delta_{s_i}) g_k(\sum_{i=1}^n \delta_{\tau(s_i)}) \mu'(ds_1) \prod_{i=2}^n \mu''(ds_i)$$

is also continuous. If $N > 0$ is any bound on u , then, for all μ' and μ'' in \mathcal{M} ,

$$\begin{aligned} |G(\mu', \mu'') - G_k(\mu', \mu'')| &\leq \sum_{n=1}^{\infty} q_n \int \int \cdots \int N |g(\sum_{i=1}^n \delta_{\tau(s_i)}) - g_k(\sum_{i=1}^n \delta_{\tau(s_i)})| \mu'(ds_1) \prod_{i=2}^n \mu''(ds_i) \\ &= N \sum_{n=0}^{\infty} \frac{1}{n!} e^{-\mathbb{E}\mathbf{P}(\mathcal{T})} \int \int \cdots \int n |g(\sum_{i=1}^n \delta_{t_i}) - g_k(\sum_{i=1}^n \delta_{t_i})| \prod_{i=1}^n \mathbb{E}\mathbf{P}(dt_i) \\ &= N [\mathbf{P}(\mathcal{T}) |g(\mathbf{P}) - g_k(\mathbf{P})|] < N/k. \end{aligned}$$

Thus, $G_k \rightarrow G$ uniformly in $\mathcal{M} \times \mathcal{M}$, and the restriction of G to this set is therefore continuous. This proves that the symmetric two-person game defined above has a symmetric equilibrium, (μ, μ) . As shown, every mixed-strategy profile $t \mapsto \mu_t$ that represents the distributional strategy μ is a mixed-strategy equilibrium in \mathbf{T} . ■

Proof of Proposition 3. Suppose that the type space \mathcal{T} is countable and all strategy sets are finite. It has to be shown that every random-player game \mathbf{T} has a mixed-strategy equilibrium. Note that the definitions of mixed-strategy profile and equilibrium do not refer to the topology on the space of strategies \mathcal{S} , but only to the measurable structure it generates. If \mathcal{S} is given the discrete topology, then \mathcal{T} and $\mathcal{S} \times \mathcal{N}_{\mathcal{S}}$ are also discrete. A subset K of \mathcal{T} is then compact if and only if it is finite, and in this case, $\tau^{-1}(K)$ is also finite, since all strategy sets are finite. Any function on a discrete topological space is continuous, and any finite point process on the (countable) space \mathcal{T} is an extended sample process. Therefore, by Theorem 3, there is a mixed-strategy equilibrium in \mathbf{T} . ■

Proof of Theorem 4. Each set F_i is both open and closed. Therefore, if $(s_k, S_k) \rightarrow (s, S)$ is a converging sequence in $\mathcal{S} \times \mathcal{N}_{\mathcal{S}}$, then $S_k(F_i) = S(F_i)$ for all i and all k large enough, and hence $u(s_k, S_k) \rightarrow u(s, S)$. Thus, u is continuous. It follows, by Theorem 3, that if \mathbf{T} is a simple mixed sample process, then it has a mixed-strategy equilibrium, $t \mapsto \mu_t$. Let \mathbf{S} be the correlated equilibrium obtained from \mathbf{T} through $t \mapsto \mu_t$. By Theorem 1 and Proposition 2, \mathbf{T}_t , and hence also \mathbf{S}_t , are the same for all $t \in \mathcal{T}$. Both are mixed sample processes, with the same sample-size distribution, and the latter is obtained from the former through the transition probability $t \mapsto \mu_t$. In \mathbf{T}_t , the probability measure according to which the elements in the sample are drawn is the same as in \mathbf{T} , and in \mathbf{S}_t , the same as in \mathbf{S} .

For every t , $\sum_{i=1}^m \mu_t(F_i) = 1$. Therefore, \mathcal{T} can be partitioned into m measurable sets B_1, B_2, \dots, B_m such that $\mu_t(F_{i(t)}) > 0$ for all t , where $i(t)$ is defined by the relation $t \in B_{i(t)}$. If $\mathbb{E}T$ is nonatomic, then these sets can moreover be chosen in such a way that $\mathbb{E}T(B_i) = \int \mu_t(F_i) \mathbb{E}T(dt)$ for all i (Dvoretzky *et al.*, 1951), and hence $\mathbb{E}T(B_i) = \mathbb{E}S(F_i)$ by Eq. (5). The set $\{(s, t) \in \mathcal{S} \times \mathcal{T} \mid s \in F_{i(t)} \text{ and } \tau(s) = t\}$ is measurable, and its \mathcal{T} -sections are compact and nonempty. The function $(s, t) \mapsto \mathbb{E}[u(s, S_t)]$ is measurable, and is continuous in s for every $t \in \mathcal{T}$. Since \mathcal{S} and \mathcal{T} are Polish spaces, these facts guarantee the existence of a measurable function $\sigma: \mathcal{T} \rightarrow \mathcal{S}$ such that $\sigma(t) \in F_{i(t)}$, $\tau(\sigma(t)) = t$, and $\mathbb{E}[u(\sigma(t), S_t)] = \max_{s \in F_{i(t)} \cap \tau^{-1}(\{t\})} \mathbb{E}[u(s, S_t)]$ for all t (Wagner, 1977, Theorem 9.2). Since $t \mapsto \mu_t$ is a mixed-strategy equilibrium and $\mu_t(F_{i(t)}) > 0$ for all t , $\max_{s \in \tau^{-1}(\{t\})} \mathbb{E}[u(s, S_t)] = \max_{s \in F_{i(t)} \cap \tau^{-1}(\{t\})} \mathbb{E}[u(s, S_t)]$ for $\mathbb{E}T$ -almost all t . Therefore, to prove that σ is a pure-strategy equilibrium it suffices to show that, for every s , $\mathbb{E}[u(s, S_t)] = \mathbb{E}[u(s, T_t \circ \sigma^{-1})]$. A sufficient condition for this is that the joint distribution of $(T_t \circ \sigma^{-1})(F_1), (T_t \circ \sigma^{-1})(F_2), \dots, (T_t \circ \sigma^{-1})(F_m)$ is equal to that of $S_t(F_1), S_t(F_2), \dots, S_t(F_m)$. Since T_t and S_t are mixed sample processes that differ only in the probability measures according to which the elements in the sample are drawn, which are the same as in T and S , respectively, the joint distributions are equal if and only if (i) $T_t(\mathcal{T}) = 0$ almost surely or (ii) $\mathbb{E}T(\sigma^{-1}(F_i)) = \mathbb{E}S(F_i)$ for all i . If $T(\mathcal{T}) < 2$ almost surely, then (i) holds. If not, then $\mathbb{E}T$ is nonatomic (see Section 2.1), and hence, by the way the B_i 's were chosen, $\mathbb{E}T(B_i) = \mathbb{E}S(F_i)$ for all i . This implies that (ii) holds. Indeed, since, for every i and every $t \in B_i$, the relation $\sigma(t) \in F_i$ holds, the equality $\sigma^{-1}(F_i) = B_i$, and hence also $\mathbb{E}T(\sigma^{-1}(F_i)) = \mathbb{E}T(B_i)$, hold for all i . ■

The following characterization of correlated equilibria is used in the proof of Theorem 5. It shows, among other things, that, under the continuity assumption of Lemma 1, the set of all distributions of correlated equilibria is a convex measurable subset of $\Delta(\mathcal{N}_S)$.

LEMMA 5. *Suppose that, for every play S and every player type t , the restriction of $u(\cdot, S)$ to the strategy set of t is continuous. Then, there is a sequence $\{\sigma_m\}_{m \geq 1}$ of pure-strategy profiles and a sequence $\{f_k\}_{k \geq 1}$ of continuous functions from \mathcal{S} to the unit interval $[0, 1]$ such that a necessary and sufficient condition for a correlated strategy S to be a correlated equilibrium is that*

$$(15) \quad \mathbb{E}\left[\int u(s, S - \delta_s) f_k(s) S(ds)\right] \geq \mathbb{E}\left[\int u(\sigma_m(\tau(s)), S - \delta_s) f_k(s) S(ds)\right]$$

for all m and k . If this condition holds, then a similar inequality to (15) holds with σ_m and f_k replaced by any pure-strategy profile and bounded continuous nonnegative-valued function, respectively.

Proof. The strategy set $\tau^{-1}(\{t\})$ of every player type t is closed and nonempty, and the multivalued function $t \mapsto \tau^{-1}(\{t\})$ is upper semicontinuous (see footnote 2). Since \mathcal{S} is a Polish space, these facts guarantee the existence of a sequence $\{\sigma_m\}_{m \geq 1}$ of measurable functions from \mathcal{T} to \mathcal{S} such that, for every t , $\{\sigma_m(t)\}_{m \geq 1}$ is a dense subset of $\tau^{-1}(\{t\})$ (Wagner,

1977, Theorem 4.2). It follows from the continuity assumption of the lemma that, for every correlated strategy \mathcal{S} and every strategy s , $\mathbb{E}[u(s, \mathcal{S}_s)] = \sup_{s' \in \tau^{-1}(\{\tau(s)\})} \mathbb{E}[u(s', \mathcal{S}_s)]$ if and only if $\mathbb{E}[u(s, \mathcal{S}_s)] \geq \mathbb{E}[u(\sigma_m(\tau(s)), \mathcal{S}_s)]$ for all m . Therefore, \mathcal{S} is a correlated equilibrium if and only if, for every m , $\mathbb{E}[u(s, \mathcal{S}_s) - u(\sigma_m(\tau(s)), \mathcal{S}_s)] \geq 0$ for $\mathbb{E}\mathcal{S}$ -almost all s . Since \mathcal{S} is a metrizable space, this condition holds if and only if, for all m and all continuous functions $f: \mathcal{S} \rightarrow [0, 1]$,

$$(16) \quad \int \mathbb{E}[u(s, \mathcal{S}_s) - u(\sigma_m(\tau(s)), \mathcal{S}_s)] f(s) \mathbb{E}\mathcal{S}(ds) \geq 0$$

(Dunford and Schwartz, 1958, Theorem IV.6.2). Clearly, if σ_m is replaced by any other pure-strategy profile, (16) is still a *necessary* condition for \mathcal{S} to be a correlated equilibrium.

By Lemma 4, the set of all mean measures of correlated strategies is compact, and hence tight. Therefore, there is an increasing sequence K_1, K_2, \dots of compact sets in \mathcal{S} such that these measures are all supported in $\bigcup_{i \geq 1} K_i$. The set of all continuous functions from a compact metric space to the unit interval is separable in the topology of uniform convergence. Therefore, by the Tietze extension theorem, there exists a sequence of continuous functions $f_k: \mathcal{S} \rightarrow [0, 1]$ ($k = 1, 2, \dots$) such that, for every continuous function $f: \mathcal{S} \rightarrow [0, 1]$, there is a subsequence $\{f_{k_i}\}_{i \geq 1}$ that converges to f uniformly on each of the K_i 's. It follows that (16) holds for every such f if and only if $\int \mathbb{E}[u(s, \mathcal{S}_s) - u(\sigma_m(\tau(s)), \mathcal{S}_s)] f_k(s) \mathbb{E}\mathcal{S}(ds) \geq 0$ for all k . By (2), this inequality is equivalent to (15). ■

Proof of Theorem 5. Let T be a finite point measure on \mathcal{T} , and B its (finite) support. Obviously, T can be seen as a finite point measure on B . Consider the random-player game in which the space of strategies is $\tau^{-1}(B)$, the payoff function and the type map are the relevant restrictions of u and τ , and the types of the players are given (deterministically) by T . The continuity assumption of the theorem implies that the payoff function in this “restricted” game is continuous. Therefore, by Theorem 3, it has at least one mixed-strategy equilibrium. Let $\{\sigma_m\}_{m \geq 1}$ and $\{f_k\}_{k \geq 1}$ be as in Lemma 5. By the second part of this lemma, the existence of mixed-strategy, and hence also correlated, equilibrium in the restricted game implies that there is a finite point process \mathcal{S} on \mathcal{S} such that (i) $\mathcal{S} \circ \tau^{-1} = T$ almost surely and (ii) the inequality (15) holds for all m and k .

Let Δ be the set of all finite point processes \mathcal{S} with finite mean measure on \mathcal{S} that satisfy (ii). Since both sides of (15) define measurable functions on $\Delta(\mathcal{N}_{\mathcal{S}})$, this set is measurable. For $T \in \mathcal{N}_{\mathcal{T}}$, let $\psi(T)$ be the set of all distributions of finite point processes \mathcal{S} on \mathcal{S} that satisfy (i). By Lemma 4, this set is compact. Using the Skorohod’s representation theorem (Billingsley, 1999, Theorem 6.7), for example, it is not difficult to show that the (nonempty) intersection $\psi(T) \cap \Delta$ is a closed subset of $\Delta(\mathcal{N}_{\mathcal{S}})$, and is therefore compact too. The graph of the multivalued function $T \mapsto \psi(T) \cap \Delta$ is a measurable subset of $\mathcal{N}_{\mathcal{T}} \times \Delta(\mathcal{N}_{\mathcal{S}})$, since the graph of $T \mapsto \psi(T)$ is easily seen to be closed and the set Δ is measurable. Since $\mathcal{N}_{\mathcal{T}}$ and $\Delta(\mathcal{N}_{\mathcal{S}})$ are Polish spaces, these facts guarantee the existence of a measurable function from $\mathcal{N}_{\mathcal{T}}$ to $\Delta(\mathcal{N}_{\mathcal{S}})$ that assigns each finite point measure T an element of $\psi(T) \cap \Delta$ (Wagner, 1980, Theorem 3.5). Therefore, for every $T \in \mathcal{N}_{\mathcal{T}}$ there is a finite point process \mathcal{S}_T on \mathcal{S} with $\mathcal{S}_T \circ \tau^{-1} = T$ almost surely and $\mathcal{S}_T \in \Delta$, such that the function $T \mapsto \mathcal{S}_T$ is measurable. Mixing \mathcal{S}_T with

respect to the distribution of \mathbf{T} gives a finite point process \mathbf{S} on \mathcal{S} that satisfies (4). In particular,

$$\mathbb{E}[g(\mathbf{S} \circ \tau^{-1})] = \mathbb{E}_{\mathbf{T}} [\mathbb{E}_{\mathbf{S}_{\mathbf{T}}} [g(\mathbf{S}_{\mathbf{T}} \circ \tau^{-1})]] = \mathbb{E}[g(\mathbf{T})] \quad (g \in \mathfrak{F}(\mathcal{N}_{\mathcal{T}})),$$

which shows that \mathbf{S} is a correlated strategy. Also,

$$\mathbb{E}\left[\int h(s, \mathbf{S} - \delta_s) \mathbf{S}(ds)\right] = \mathbb{E}_{\mathbf{T}} \left[\mathbb{E}_{\mathbf{S}_{\mathbf{T}}} \left[\int h(s, \mathbf{S}_{\mathbf{T}} - \delta_s) \mathbf{S}_{\mathbf{T}}(ds)\right]\right] \quad (h \in \mathfrak{F}(\mathcal{S} \times \mathcal{N}_{\mathcal{S}})),$$

and since $\mathbf{S}_{\mathbf{T}} \in \Delta$ for all \mathbf{T} , this equality implies that $\mathbf{S} \in \Delta$, too. Therefore, by Lemma 5, \mathbf{S} is a correlated equilibrium. ■

Proof of Proposition 4. First, it has to be shown that if \mathbf{S} is a correlated strategy in the first random-player game (i.e., $\mathbf{S} \circ \tau^{-1} = \mathbf{T}$), then $(\mathbf{S} \circ \eta^{-1}) \circ \hat{\tau}^{-1} = \hat{\mathbf{T}}$. Since, by condition (iii), $\hat{\mathbf{T}} = \mathbf{T} \circ \theta^{-1} = (\mathbf{S} \circ \tau^{-1}) \circ \theta^{-1} = \mathbf{S} \circ (\theta \circ \tau)^{-1}$, it suffices to show that $\theta \circ \tau = \hat{\tau} \circ \eta$. Let s, t , and \hat{t} be such that $s \in \tau^{-1}(\{t\})$ and $\theta(t) = \hat{t}$. Then, $\theta(\tau(s)) = \hat{t}$. On the other hand, by condition (ii), $\eta(s) \in \tau^{-1}(\{t\})$, and hence $\hat{\tau}(\eta(s)) = \hat{t}$. Therefore, $\theta \circ \tau = \hat{\tau} \circ \eta$, as had to be shown.

Next, suppose that \mathbf{S} is a correlated equilibrium. Let $\hat{\psi}: \hat{\mathcal{S}} \rightarrow \hat{\mathcal{S}}$ be some measurable function such that $\hat{\tau} \circ \hat{\psi} = \hat{\tau}$. If there is a measurable function $\phi: \mathcal{S} \rightarrow \mathcal{S}$ such that $\tau \circ \phi = \tau$ and $\hat{\psi} \circ \eta = \eta \circ \phi$, then

$$\begin{aligned} \mathbb{E}\left[\int \hat{u}(\hat{\psi}(\hat{s}), \mathbf{S} \circ \eta^{-1})(\mathbf{S} \circ \eta^{-1})(d\hat{s})\right] &= \mathbb{E}\left[\int \hat{u}((\hat{\psi} \circ \eta)(s), \mathbf{S} \circ \eta^{-1}) \mathbf{S}(ds)\right] \\ &= \mathbb{E}\left[\int \hat{u}((\eta \circ \phi)(s), \mathbf{S} \circ \eta^{-1}) \mathbf{S}(ds)\right] = \mathbb{E}\left[\int u(\phi(s), \mathbf{S}) \mathbf{S}(ds)\right] \\ &\leq \mathbb{E}\left[\int u(s, \mathbf{S}) \mathbf{S}(ds)\right] = \mathbb{E}\left[\int \hat{u}(\hat{s}, \mathbf{S} \circ \eta^{-1})(\mathbf{S} \circ \eta^{-1})(d\hat{s})\right], \end{aligned}$$

where the third equality follows from condition (i) and the inequality from Lemma 1. Therefore, again by Lemma 1, to prove that $\mathbf{S} \circ \eta^{-1}$ is a correlated equilibrium it suffices to show that for every $\hat{\psi}$ as above there is a measurable function $\phi: \mathcal{S} \rightarrow \mathcal{S}$ that renders the following diagram commutative:

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{(\tau, \hat{\psi} \circ \eta)} & \mathcal{T} \times \hat{\mathcal{S}} \\ & \searrow \phi & \nearrow (\tau, \eta) \\ & \mathcal{S} & \end{array}$$

For every $s \in \mathcal{S}$, $\hat{\tau}(\hat{\psi}(\eta(s))) = (\hat{\tau} \circ \hat{\psi})(\eta(s)) = \hat{\tau}(\eta(s))$, and it therefore follows from condition (ii) that $\hat{\psi}(\eta(s)) = \eta(s')$ for some $s' \in \mathcal{S}$ such that $\tau(s) = \tau(s')$. Thus, the range of the function $(\tau, \hat{\psi} \circ \eta)$ is contained in that of (τ, η) . It follows from the continuity of η and the properties of the type map τ that, for every s , the set of all strategies s' as above is compact. The set of all pairs $(s, s') \in \mathcal{S} \times \mathcal{S}$ with elements related as above is measurable. Since \mathcal{S} is a Polish space, these facts guarantee the existence of a measurable function $\phi: \mathcal{S} \rightarrow \mathcal{S}$ such that, for every s , $(s, \phi(s))$ belongs to the last-mentioned set (Wagner, 1980, Theorem 3.5). This function ϕ renders the above diagram commutative. ■

Proof of Proposition 5. Given the random-player game \hat{T} and the correlated strategy \hat{S} , the random-player game T is defined as follows. The space of strategies is defined by $\mathcal{S} = \{(\hat{s}, \hat{s}') \in \hat{\mathcal{S}} \times \hat{\mathcal{S}} \mid \hat{\tau}(\hat{s}) = \hat{\tau}(\hat{s}')\}$. The payoff function is given by $u(s, S) = \hat{u}(\pi_2(s), S \circ \pi_2^{-1})$, where $\pi_2 : \mathcal{S} \rightarrow \hat{\mathcal{S}}$ is the projection on the second coordinate. The type map τ is given by $(\hat{s}, \hat{s}') \mapsto (\hat{s}, \hat{\tau}(\hat{s}))$, and the type space \mathcal{T} thus coincides with the graph of $\hat{\tau}$ (which entails that player types in the random-player game T essentially correspond to strategies in \hat{T}). Finally, $T = \hat{S} \circ (\text{id}, \hat{\tau})^{-1}$, where id is the identity map on $\hat{\mathcal{S}}$. It is easy to see (e.g., by using Lemma 1) that the correlated strategy $S = \hat{S} \circ (\text{id}, \text{id})^{-1}$ is a correlated equilibrium in this random-player game. This correlated equilibrium represents the pure-strategy profile $(\hat{s}, \hat{t}) \mapsto (\hat{s}, \hat{s})$, which is therefore a pure-strategy equilibrium. The function $\eta = \pi_2$ transforms S into \hat{S} (i.e., $\hat{S} = S \circ \eta^{-1}$). This function is easily seen to satisfy conditions (i)–(iii). ■

REFERENCES

- Aumann, R. J., 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* **55**, 1–18.
- Billingsley, P., 1999. *Convergence of Probability Measures*, 2nd Edition. Wiley, New York.
- Brandenburger, A., Dekel, E., 1993. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory* **59**, 189–198.
- Cotter, K. D., 1991. Correlated equilibrium in games with type-dependent strategies. *Journal of Economic Theory* **54**, 48–68.
- Cotter, K. D., 1994. Type correlated equilibria for games with payoff uncertainty. *Economic Theory* **4**, 617–627.
- Daley, D. J., Vere-Jones, D., 1988. *An Introduction to the Theory of Point Processes*. Springer-Verlag, New York.
- Dunford, N., Schwartz, J. T., 1958. *Linear Operators. Part I*. Interscience, New York.
- Dvoretzky, A., Wald, A., and Wolfowitz, J., 1951. Relations among certain ranges of vector measures. *Pacific Journal of Mathematics* **1**, 59–74.
- Fan, K., 1952. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences of the USA* **38**, 121–126.
- Forges, F., 1993. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision* **35**, 277–310.
- Harsanyi, J. C., 1967. Games with incomplete information played by ‘Bayesian’ players. Part I. The basic model. *Management Science* **14**, 159–182.
- Harstad, R. M., Kagel, J. H., and Levin, D., 1990. Equilibrium bid functions for auctions with an uncertain number of bidders. *Economics Letters* **33**, 35–40.

- Kallenberg, O., 1986. *Random Measures*, 4th Edition. Akademie-Verlag, Berlin, Academic Press, London.
- Mailath, G. J., Samuelson, L., and Shaked, A., 1997. Correlated equilibria and local interactions. *Economic Theory* **9**, 551–556.
- Matthews, S., 1987. Comparing auctions for risk averse buyers: a buyer's point of view. *Econometrica* **55**, 633–646.
- McAfee, R. P., McMillan, J., 1987. Auctions with a stochastic number of bidders. *Journal of Economic Theory* **43**, 1–19.
- Mertens, J.-F., Zamir, S., 1985. Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory* **14**, 1–29.
- Milgrom, P. R., Weber, R. J., 1985. Distributional strategies for games with incomplete information. *Mathematics of Operations Research* **10**, 619–632.
- Myerson, R. B., 1991. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA.
- Myerson, R. B., 1998a. Population uncertainty and Poisson games. *International Journal of Game Theory* **27**, 375–392.
- Myerson, R. B., 1998b. Extended Poisson games and the Condorcet jury theorem. *Games and Economic Behavior* **25**, 111–131.
- Myerson, R. B., 2000. Large Poisson games. *Journal of Economic Theory* **94**, 7–45.
- Parthasarathy, K. R., 1967. *Probability Measures on Metric Spaces*. Academic Press, New York.
- Resnick, S. I., 1987. *Extreme Values, Regular Variation, and Point Processes*. Springer-Verlag, New York.
- Wagner, D. H., 1977. Survey of measurable selection theorems. *SIAM Journal on Control and Optimization* **15**, 859–903.
- Wagner, D. H., 1980. Survey of measurable selection theorems: an update. In: Kozlov, D. (Ed.), *Measure Theory, Oberwolfach 1979*, Lecture Notes in Mathematics, 794. Springer-Verlag, Berlin, pp. 176–219.