EXTREMAL CHOICE EQUILIBRIUM WITH APPLICATIONS TO LARGE GAMES, STOCHASTIC GAMES, & ENDOGENOUS INSTITUTIONS

PAULO BARELLI AND JOHN DUGGAN

Abstract. We prove existence and purification results for strategic environments possessing a product structure that includes classes of large games, stochastic games, and models of endogenous institutions. Applied to large games, the results yield existence of pure-strategy equilibria allowing for infinite-dimensional externalities. Applied to stochastic games, the results yield existence of stationary Markov perfect equilibria with extremal payoffs, which in turn yields existence of pure strategy stationary Markov perfect equilibria for games with sequential moves. Applied to the model of institutions, we obtain equilibrium existence with general group decision correspondences.

1. Introduction

We study the question of existence of solutions in a general class of strategic environments that includes large games, stochastic games, and endogenous institutions. In the context of large games, we obtain existence of Nash equilibria, and when underlying pure action sets are finite, our results deliver existence in pure strategies. Ours is the first such result that accommodates infinite-dimensional externalities but still uses standard measure spaces (e.g., $[0, 1]$ with Lebesgue measure), as opposed to saturated measure spaces. For stochastic games, we extend the existence theorem of Duggan (2012a) for noisy stochastic games to obtain stationary Markov perfect equilibria such that payoffs are in the closure of the extreme points of the set of Nash equilibrium

Date: November 25, 2014.

Barelli: Dept. of Economics, University of Rochester, paulo.barelli@rochester.edu. Duggan: Dept. of Political Science and Dept. of Economics, University of Rochester, dugg@ur.rochester.edu. We thank the editor and an anonymous referee for helpful comments and suggestions.

payoffs of the associated auxiliary one-shot games. In the special case of sequential move games, our results imply existence of stationary Markov perfect equilibrium in pure strategies. Finally, we consider a framework of endogenous institutions such that a large number of individuals sort into groups and then take collective decisions within groups. Generalizing a result due to Caplin and Nalebuff (1997), we establish existence of an equilibrium when group decisions can be multi-valued, as is the case for common voting rules.

We obtain these results as applications of an abstract existence theorem in a general setting characterized by a product structure of the form $T \times U$, where the interpretation of the spaces T and U depends on the application considered. In the large game application, we view $T \times U$ as the set of players, so we identify a player with a pair (t, u), where t is a general characteristic and u is a personal characteristic. We assume payoffs depend on own actions and the profile of average actions across general characteristics. That is, we "integrate out" personal characteristics, and we let average actions vary arbitrarily across general characteristics to accommodate infinite-dimensional externalities. In the stochastic games application, we view $T \times U$ as the set of states, where t is a general component of the state and u is a noise component that is payoff relevant in the current period but is not directly affected by last period's state and actions. Applied to endogenous institutions, we view $T \times U$ as a society consisting of individuals who must select into groups, where t is a public characteristic of an individual that affects others and u is a private characteristic.[1] In particular, our arguments permit the space T to be infinite by exploiting this product structure, integrating out the idiosyncratic variable u and using known results on parameterized integrals of correspondences to verify key continuity conditions. Adding an assumption of non-atomicity of the idiosyncratic variable u, we sharpen our results to deliver solutions with extremal properties, which translate into pure strategy existence results in applications.

Our main existence theorem may be viewed as an abstract fixed point result, with no immediate interpretation in terms of a game—there are, for example,

[1]For a further application, in Barelli and Duggan (2014), we show that if the product structure is imposed on type spaces in a Bayesian game, and types consist of a general component together with an atomless, conditionally independent, private-value component, then existence and purification results hold for games with finitely many actions.

no players and no payoffs functions—but it exploits structure that appears in a number of economic environments, as evidenced by the applications we provide. To illustrate the usefulness of the result, consider the following example of a static model of competition among firms.

Cournot Game among Many Firms. Consider a market composed of many firms, where each firm is characterized by its location t and technological characteristic u. Assume that there is a continuum of locations and of technologies, each represented by the unit interval $[0,1]$. Each firm $i = (t,u)$ produces a vector $q(i) \in \Re^d$ of commodities belonging to a production set $Q(i)$. Let $\alpha(t) = \int_u q(t,u)du$ denote the aggregate production vector at location t, averaging over technologies u. Assume that prices are determined by product and factor market clearing in each location, where consumers and workers may (at some cost) travel to transact in markets at different locations, so that prices and profits depend on the aggregate production function α. We can then write the profit of firm i producing vector q given aggregate production α as $\pi_i(q;\alpha)$. Thus, we have a large game with $N = T \times U$ being the player set (the firms), where $T = U = [0,1]$ are the spaces of location and technologies (endowed with the standard Lebesgue measure), $Q(i)$ being the action set, and π_i being the payoff function of player $i \in N$. The externality caused by other firms on a given firm i is summarized by the infinite-dimensional aggregate production function α. This is to be contrasted with similar models in the literature that either only allow for finite-dimensional externalities (see, e.g., Wu and Zhu (2005)) or manage to allow for infinite-dimensional externalities by imposing a much richer measurable structure on T and U (i.e., by using a super-atomless measure instead of the Lebesgue measure, as in Carmona and Podczeck (2014).) Let $M(i;\alpha)$ denote the best response correspondence of firm i: that is, given the aggregate production α determined by the choices of all other firms, we compute the set of production vectors $q \in Q(i)$ that maximize $\pi_i(q;\alpha)$. An equilibrium of the game is thus a function $q^*(\cdot)$ such that $q^*(i) \in M(i;\alpha^*)$ for all firms i, where α^* is the aggregate production associated with $q^*(\cdot)$, i.e., $\alpha^*(t) \equiv \int_u q^*(t,u)du$. As we argue next, our general fixed point result allows us to establish fairly general conditions ensuring existence (and purification) of equilibria in such a general Cournot game.

Analytical approach. Our general framework is formulated abstractly, and our main theorem can be viewed as a fixed point theorem that exploits a product structure on its domain.[2] To convey the idea, we define a choice function γ as assigning to each pair (t,u) a choice in $\mathfrak{R}^d$, where we view t as a systematic variable and u as an idiosyncratic variable. We then calculate the corresponding average choice function, α, by taking the marginal, $\alpha(t) = \int_u \gamma(t,u)du$, of γ across the idiosyncratic variable pointwise for each t. We then assign a choice set $M(t,u;\alpha)$ to each pair (t,u), where by construction these sets are parameterized by average choices, and we define a choice equilibrium as a mapping γ such that for all (t,u) pairs, $\gamma(t,u)$ belongs to the choice set $M(t,u;\alpha)$ determined by the corresponding average choices. The firm competition example above is immediately seen as a special case of the general framework upon setting $\gamma = q$. Beyond existence, assuming u is non-atomically distributed, we provide an "extremization" result: for every choice equilibrium, there is an extremal choice equilibrium $\hat{\gamma}$ that chooses from the (closure of) extreme points of choice sets $M(t,u;\hat{\alpha})$ such that $\hat{\gamma}$ is equivalent to γ, in the sense that it determines the same average choices and, therefore, the same choice sets for all (t,u) pairs.

The existence argument takes place in the space of average choice functions. We define $S(\alpha)$ as the set of selections of the correspondence $t \mapsto \int_u M(t,u;\alpha)du$, and we prove existence of a fixed point $\alpha^* \in S(\alpha^*)$ that is generated by an equilibrium choice function γ^*. The fixed point argument surmounts a number of technical challenges. To ensure sequential upper hemicontinuity of S, we apply a result of Artstein (1979) on weak limits of sequences of integrable functions, and as the space of average choice functions is not necessarily (weakly) compact or metrizable, we apply a recent result of Agarwal and O'Regan (2002) to obtain a fixed point, α^*. Finally, we employ the theorem of Artstein (1989) to back out an equilibrium choice function γ^* consistent with α^*. Our purification argument relies on an application of a version of Lyapunov's theorem pointwise for each t, using non-atomicity of u; we then apply Artstein's theorem again to back out an extremal choice function. The latter step relies on a result establishing lower measurability of the extreme points of a lower measurable correspondence with nonempty, compact values in $\mathfrak{R}^d$.

Applications. We apply the results above to large games by indexing players as (t,u), where t is a general characteristic of the player and u is a

[2]This product structure allows us to apply the iterated integral approach used by Duggan (2012a) to prove existence of stationary Markov perfect equilibria in noisy stochastic games.

personal characteristic, and we interpret choice set $M(t,u;\alpha)$ as the set of best response actions of player (t,u), given the average action α as a function of the general characteristic. Thus, given a strategy profile σ, the average action $\alpha(t) \equiv \int_u \sigma(t,u)du$ is an infinite-dimensional statistic of the actions of the players. We then consider a large game such that this statistic (the "externality") captures the influence of the rest of the players on a given player's outcome, and a choice equilibrium corresponds to a Nash equilibrium of the large game. We provide general sufficient conditions on the players' preferences to generate well-behaved choice sets $M(t,u;\alpha)$ and therefore existence of an extremal equilibrium. We illustrate the approach with the above-mentioned example of a Cournot game with many firms characterized by their location and production technology: our results deliver existence of equilibrium while allowing market clearing conditions in infinitely many locations to affect to price received by a given firm, generalizing results in the literature. When feasible action sets have a simplicial structure and payoffs are multilinear, our result delivers existence in pure strategies, the first such result using standard measure spaces; the cost is that underlying pure action sets must be finite, whereas with saturated measure spaces they can be more general.

The application to noisy stochastic games interprets a pair (t,u) as a decomposition of a state in a stochastic game, with the first component t satisfying the usual assumptions in the literature and the second component u being conditionally independent of the previous state and actions. That is, t is a general component of the state, and u is the "noise" component, as in Duggan (2012a). Choice sets $M(t,u;\alpha)$ are the sets of Nash equilibrium payoffs of auxiliary one-shot games indexed by the vector of (interim) continuation values, given by the average choice α, so a choice equilibrium selects Nash equilibria from each such one-shot game. Dynamic programming ideas then show that a choice equilibrium corresponds to a stationary Markov perfect equilibrium of the noisy stochastic game. We therefore establish existence of a particular kind of stationary equilibrium, namely one that selects extreme points from the set of Nash equilibrium payoffs of auxiliary one-shot games. This refinement implies existence of a pure strategy stationary Markov perfect equilibrium for noisy stochastic games with sequential moves. We illustrate with a dynamic sequential oligopoly game with random movers, where our results guarantee existence of a pure strategy stationary Markov perfect equilibrium.

In the application to endogenous institutions, we identify an individual with a pair (t, u), where t is a public characteristic and u is a private characteristic that does not affect other individuals. We establish existence of a pure strategy equilibrium in a general framework for endogenous sorting into groups and collective decisions within groups, allowing for general spaces of individual characteristics and preferences (including crowding effects within groups), constraints on group membership that depend on individual characteristics, and for group decision correspondences, which arise naturally for common voting rules; the latter assumption generalizes Caplin and Nalebuff (1997), who restrict attention to single-valued mappings. The existence argument proceeds by transforming the model into our abstract framework by adding artificial agents to represent collective decisions within groups. Here, $M(t, u; \alpha)$ consists of mixtures over the optimal groups for the original individuals (t, u), given the distribution of general characteristics across groups summarized by α, and for the artificial agent representing group j, $M(t, u; \alpha)$ consists of the set of group decisions for the group given membership profile α. We apply our general existence and extremization results to obtain the desired equilibrium in the institutional model. We illustrate with an example of super majority voting over local public goods, in which group decisions are made by voting (according to a quota rule) and individuals select into groups, where our results guarantee existence of an endogenous institutional equilibrium.

Related Literature. Our existence result for choice equilibria in the abstract framework is comparable but non-nested with Theorem 2.2.1 of Balder (2002). The latter paper establishes existence of pure strategy equilibria in pseudogames that are more general than our framework in that his action sets may be infinite-dimensional, but less general in that it assumes externalities are finite-dimensional. At a finer level of detail, the papers also differ in how convexity conditions are formulated; see Remark 3.6 for an explanation of how the latter distinction arises from our general treatment of externalities. Note that when the idiosyncratic variable u is non-atomically distributed for all t, as would be the case in most applications, convexity is not needed in either framework, so the distinction is moot in this case.

In large games, Schmeidler (1973) first provides conditions for existence of Nash equilibrium in pure strategies with finite sets of pure actions using the integral of the strategy as the "societal response" (also called the "externality"). Since action sets are finite, such externalities are finite-dimensional; our Corollary 3.2 generalizes that result by allowing richer, infinite-dimensional, externalities. Mas-Colell (1984) uses the distribution of the strategy as the externality, and under such alternative formulation is able to establish existence of mixed strategy equilibria for general action sets.[3] As indicated by Khan, Rath, and Sun (1997), once infinite action sets are considered, it makes a difference how one considers the effect of the other players' actions on a given player's payoffs. We follow their paper and others in capturing externalities as an integral (an average) rather than a distribution, and we refer to their arguments in favor of the integral approach over the distribution approach. In this class of games, our existence results for Nash equilibria in large games is non-nested with respect to the results in Martins-da-Rocha and Topuzu (2008) and Balder (2002), as we allow for infinite-dimensional externalities at the cost of finite-dimensional action sets. With respect to Khan, Rath, and Sun (1997), we provide a modeling approach that allows us to handle infinite-dimensional externalities (as integrals) without relying on an infinite-dimensional version of Lyapunov's theorem. In particular, letting σ denote a strategy profile and $\sigma(t,u)$ denote the action of player (t,u), the approach in Khan, Rath, and Sun (1997) would be to condense externalities to the finite-dimensional statistic $\beta = \int_{(t,u)} \sigma(t,u)d(t,u)$, with the implication that two strategy profiles σ and $\hat{\sigma}$ with $\beta = \hat{\beta}$ would be considered equivalent by all players. In contrast, in our model, it is the infinite-dimensional statistic $\alpha(\cdot) = \int_u \sigma(\cdot,u)du$ on which players condition their choices; it is obviously possible to have $\alpha \neq \hat{\alpha}$ while $\beta = \hat{\beta}$, so players react to a richer set of "societal statistics" in our formulation.

This modeling strategy circumvents the failure of Lyapunov's theorem in infinite dimensions, without resorting to extremely diffuse environments (i.e., saturated measure spaces with super-atomless measures), as in Podczeck (2009), Carmona and Podczeck (2009,2014), and Khan, Rath, Sun, and Yu (2013).[4]

[3] In the finite-action setting of Schmeidler (1973), the two formulations of the externality are equivalent.

[4] In the saturated/super-atomless environment, more mathematically sweeping results can be obtained. In fact, Carmona and Podczeck (2014) obtain existence of pure strategy equilibria allowing not only for general action spaces and infinite-dimensional externalities, but also for discontinuous payoff functions. As shown by Greinecker and Podczeck (2014), however, purification in such environments is spurious.

It is worth noting that, under Mas-Colell's distribution approach for a large game, Noguchi (2009) has recently proved existence of a pure strategy equilibrium allowing for general action sets, under the assumption that there are "many agents of every type." This assumption allows him to circumvent the failure of Lyapunov's theorem in infinite dimensions without resorting to extremely diffuse environments, but in our framework, the assumption means that u is payoff irrelevant; as such, it reduces to a randomization device for player t, limiting the usefulness of the results for applications.

The existence of an extremal stationary Markov perfect equilibrium in a noisy stochastic game is new, as is the extension to pure strategy equilibria in sequential move games. These results build on the recent result of Duggan (2012), which in turn generalizes the result of Nowak and Raghavan (1992). We note that the product structure on states is crucial in obtaining existence of pure strategy stationary equilibria in sequential move games: restricting to sequential move games in the framework of Nowak and Raghavan (1992) only allows one to remove correlation from their equilibrium; with the product structure, we can remove mixing altogether. The pure strategy existence result generalizes Theorem 5 in Duggan and Kalandrakis (2012), which holds for a special class of sequential move stochastic games such that states are finite-dimensional, the transition on the general component of the state is smooth, and the noise component of the state enters into payoffs in a non-degenerate way. On the other hand, they show that all equilibria are essentially pure and possess desirable continuity properties, results we do not replicate in our more general framework. Recently, Jaśkiewicz and Nowak (2014) show existence of a pure strategy Markov perfect equilibrium in stochastic overlapping generations models with sequential moves allowing for more general payoff functions (including hyperbolic discounting). Rather than assuming an atomless noise component of the state, those authors assume that the state transition is a convex combination of a fixed, finite set of atomless transition probabilities.

Models of local public goods with mobility trace back to early work of Tiebout (1956), and existence of equilibrium has been a focus since Westhoff (1997), Epple et al. (1984), and Konishi (1996). We adopt the general framework of Caplin and Nalebuff (1997), who show that equilibria may fail to exist, due to an incentive of some types of individuals to "run toward" others, and an incentive for the latter types to "run from" the former. Those authors propose three routes to equilibrium existence, the second of which endows the set of individuals with a product structure, so that an individual is identified as (t, u), where t is a general characteristic that may affect the well-being of others, and $u \in \mathfrak{R}^n$ is a vector of additive preference shocks on the utility of belonging to any given institution. Our results maintain the decomposition of individuals into public and private characteristics and generalize this route to existence in several ways: we let the set of individuals be the product of Polish spaces, rather than Euclidean spaces, and we weaken continuity of utility to measurability in the individuals' types; we allow the private component u to enter into individual preferences in a general way, instead of the additive form; we allow for arbitrary membership constraints that restrict entry into the groups; we allow an individual to care not only about the vector of group decisions but also about the composition of the group to which she belongs; and we allow for group decision correspondences, which associate sets of possible decisions to each group given the selection of individuals into groups.

Organization. In Section 2, we present the abstract framework and our general existence and purification result. Section 3 provides an application of our general results to large games, Section 4 takes up the case of noisy stochastic games, and Section 5 presents the application to endogenous institutions. The abstract theorem is proved in Appendix A; Appendix B contains the proofs of the results in Sections 3–5; and Appendix C contains a technical lemma on lower measurability of extreme points of a correspondence.

2. Abstract Framework and Main Result

Informal discussion. We construct an abstract framework that takes as given a set $N = T \times U$ with a product structure and a correspondence $A\colon N \rightrightarrows \mathfrak{R}^d$ such that for each $i \in N$, $A(i)$ represents a set of alternatives that are feasible at i. In the large game setting, $A(i)$ is the set of actions available to player i; in the stochastic game setting, an element $i \in N$ represents a state, and the set $A(i)$ is a fixed set that bounds the payoff vectors of the players; in the endogenous institution section, $A(i)$ is a face of the unit simplex corresponding to groups that individual i can join. The endogenous variable in our model is a selection $\gamma\colon N \to \mathfrak{R}^d$ from the correspondence A. Here, an element $i \in N$ consists of a pair $i = (t, u)$, where t is a systematic variable and u is an idiosyncratic variable. The precise meaning of these terms depends on the application intended, but roughly speaking, given a selection γ from A, it is the marginal of γ across the idiosyncratic variable u that determines a choice set of alternatives for each (t,u) pair. That is, letting $\alpha(t) = \int_u \gamma(t,u)$ denote the average choice conditional on t, choice sets are given by $M(t,u;\alpha)$, which is parameterized by the (potentially infinite-dimensional) family of average choice functions $\alpha: T \to \mathfrak{R}^d$.

We formulate an abstract equilibrium concept in terms of the inclusion $\gamma(t,u) \in M(t,u;\alpha)$ for all (t,u), where the correspondence $M(\cdot;\alpha)$ has a general form. In the large game setting, a pair (t,u) corresponds to a player, where t is a general characteristic and u is a personal characteristic of the player that does not affect the payoffs of others. Then $M(t,u;\alpha)$ is the set of best response actions for player (t,u) given average actions α. In the noisy stochastic game setting, (t,u) is a state variable, where t is a general component of the state, and u is a noise component that, while payoff relevant, is not affected by the state and actions in the previous period. Here, $\alpha(t)$ represents the vector of continuation payoffs of the players given the general state component t, and then $M(t,u;\alpha)$ is the set of mixed strategy equilibrium payoff vectors in the auxiliary one-shot game at state (t,u) induced by the continuation value α.

The formulation of endogenous institutions in our framework is slightly more complicated. In this case, we use N to represent the set of individuals (t,u), where t is a public characteristic that affects other individuals' payoffs and group decisions, and u is a private characteristic; and in addition, we fold the finite number of possible groups into T and view them as artificial players. Then for an individual (t,u), $A(t,u)$ is the set of mixtures over groups that the individual can join, and for a group t, $A(t,u)$ is equal to the space of possible collective decisions for the groups. Here, $\alpha(t)$ yields the distribution of individuals with public characteristic t across the groups, or when t represents a group, it specifies the collective decision for group t. Finally, for an individual (t,u), $M(t,u;\alpha)$ is the set of optimal group memberships for the individual, and for a group t, $M(t,u;\alpha)$ is the set of possible collective decisions for the group given α. These relationships are depicted in Figure 1.

The main result of the paper establishes existence of a choice equilibrium, which is a mapping γ from (t,u) pairs to $\mathfrak{R}^d$ such that $\gamma(t,u) \in M(t,u;\alpha)$ for all (t,u), where α is the average choice generated by γ. Without going into details at the moment, the conditions for existence are imposed on the feasible set

abstract framework	large games	noisy stochastic games	endogenous institutions
N	player set	state space	set of individuals and group names
t	general characteristic	general state component	public characteristic or group name
u	personal characteristic	noisy state component	private characteristic of individual
$A(t,u)$	feasible actions for player (t,u)	stage payoff vectors in state (t,u)	feasible groups for (t,u) and possible group decisions
$M(t,u;\alpha)$	best responses of player (t,u) to α	m.s.e. payoff vectors in auxiliary game given cont. value α	membership choices and group decisions given α

FIGURE 1. Mapping abstract framework into applications

$A(t,u)$ and the choice set $M(t,u;\alpha)$, and they are of four kinds: conditions on values of the correspondences (nonemptiness, compactness, and in some cases convexity), measurability, boundedness, and continuity. The latter condition is imposed on the choice correspondence and is especially important for the fixed point argument; specifically, we assume that for each (t,u), the set $M(t,u;\alpha)$ of chosen alternatives varies upper hemicontinuously with the average choice α, appropriately topologized. These conditions are relatively standard and, as we shall argue, are satisfied in many applications of interest.

Formal details. Proceeding to the formalities, let T and U be complete, separable metric spaces, with their respect Borel sigma-algebras, $\mathcal{T}$ and $\mathcal{U}$. Let κ be a Borel probability measure on T, and let $\nu(\cdot|\cdot){:}T\times\mathcal{U}\to[0,1]$ be a transition probability from T to U.[5] Giving $N=T\times U$ the product topology,

[5]That is, $\nu(\cdot|\cdot){:}T\times\mathcal{U}\to[0,1]$ is a mapping such that $\nu(\cdot|t)$ is a Borel probability measure on U for κ-almost all $t\in T$, $t\mapsto\mu(E|t)$ is a $\mathcal{T}$-measurable function for all $E\in\mathcal{U}$. In particular, the mapping $t\mapsto\nu(\cdot|t)$ is Borel measurable with the weak* topology on the space of Borel probability measures on U.

let $\mu = \nu(\cdot|\cdot) \otimes \kappa$ be the Borel probability measure on N induced by κ and ν.[6] For each $(t,u) \in T \times U$, let $A(t,u) \subseteq \mathfrak{R}^d$ denote a set of *feasible alternatives.* A *choice function* is a $\mathcal{T} \otimes \mathcal{U}$-measurable mapping $\gamma: T \times U \to \mathfrak{R}^d$ such that for μ-almost all (t,u), we have $\gamma(t,u) \in A(t,u)$. Let Z be a $\mathcal{T} \otimes \mathcal{U}$-measurable set containing $\{(t,u) \in T \times U : u \text{ is an atom of } \nu(\cdot|t)\}$. We assume the values of A satisfy standard conditions, and we impose a minimal measurability property on the correspondence:

(A1) for all $(t,u) \in T \times U$, $A(t,u)$ is nonempty and compact; and for each $(t,u) \in Z$, the set $A(t,u)$ is convex,

(A2) the correspondence $A: T \times U \rightrightarrows \mathfrak{R}^d$ is lower measurable,[7]

Note that convexity of feasible sets is assumed only on the set Z, which contains the atoms of ν. Under the additional assumption that the family $\{\nu(\cdot|t) : t \in T\}$ of conditional probabilities are atomless, we may set $Z = \varnothing$.[8]

Let $1 \leq p < \infty$ and q satisfying $\frac{1}{p} + \frac{1}{q} = 1$ be fixed for the remainder of the paper. We impose the following weak boundedness assumption on the correspondence of feasible alternatives:

(A3) for κ-almost all t, the mapping $u \mapsto \sup ||A(t,u)||$ is p-integrably bounded, i.e.,

$$\int_u \sup ||A(t,u)||^p \nu(du|t) \quad < \quad \infty.$$

For later use, we record the following strengthening of assumption (A3):

(A3′) the mapping $(t,u) \mapsto \sup ||A(t,u)||$ is p-integrably bounded, i.e.,

$$\int_{(t,u)} \sup ||A(t,u)||^p \mu(d(t,u)) \quad < \quad \infty.$$

[6]Specifically, for all $Q = R \times S \in \mathcal{T} \otimes \mathcal{U}$, we have $\mu(Q) = \int_R \nu(S|t)\kappa(dt)$, so that ν is a regular conditional probability mapping each t to a Borel probability measure $\nu(\cdot|t)$ on U that gives the probability $\nu(E|t)$ of each $E \in \mathcal{U}$ conditional on t.

[7]Given measurable space X and topological space Y, a correspondence $\psi: X \rightrightarrows Y$ is *lower measurable* if for all open $G \subseteq Y$, the set $\{x \in X \mid \psi(x) \cap G \neq \varnothing\}$ is measurable.

[8]One might desire to impose convexity of $A(t,u)$ only for atoms of μ, but this hope encounters a fundamental convexity problem. Consider the simplest case in which κ is uniform on $T = [0,1]$, $U = \{u\}$ is singleton, and $A(t,u) \equiv \{0,1\}$ is constant. Then μ is atomless, but we can select a sequence of selections γ^m from the feasible alternative correspondence that converges weakly to the constant function $\gamma(t,u) \equiv \frac{1}{2}$. That is, the set of selections from $A(\cdot)$ is not weakly closed.

Assumption (A3′) (and therefore (A3)) is automatically satisfied if the feasible sets are bounded by a fixed, compact subset of $\mathfrak{R}^d$, but we allow in principle for arbitrarily large alternative sets $A(i)$.

Given a choice function γ, we will be concerned with average choices as a function of the systematic variable t, as we integrate across the idiosyncratic variable u. This determines an *average choice function* denoted $\alpha: T \to \mathfrak{R}^d$, as follows: for each $t \in T$, we define the integral of γ over u by

$$\alpha(t) \quad = \quad \int_u \gamma(t,u)\nu(du|t),$$

which is Borel measurable. More precisely, given γ, define the $\mathcal{T}$-measurable function $\alpha(\cdot|\gamma): T \to \mathfrak{R}^d$ by $\alpha(t|\gamma) = \int_u \gamma(t,u)\nu(du|t)$. Then the set of average choice functions consists of any mapping that is equivalent to some $\alpha(\cdot|\gamma)$ up to a set of κ-measure zero:

$$\mathfrak{A} = \left\{ \alpha: T \to \mathfrak{R}^d : \begin{array}{c} \alpha(t) = \alpha(t|\gamma) \text{ for } \kappa\text{-almost all } t \in T \\ \text{and for some choice function } \gamma \end{array} \right\}.$$

We will sometimes suppress dependence of $\alpha(\cdot|\gamma)$ on γ without confusion in the sequel. Note that because (A3) is stated pointwise for each t, it does not imply compactness of the space of average choice functions under the weak topology.

Given each $\alpha \in \mathfrak{A}$, we assume a *choice correspondence,* denoted $M(\cdot;\alpha): T \times U \rightrightarrows \mathfrak{R}^d$, that associates a subset of feasible alternative to each pair (t,u). Assume: for each $\alpha \in \mathfrak{A}$,

- (A4) for all $(t,u) \in T \times U$, $M(t,u;\alpha) \subseteq A(t,u)$,
- (A5) for all $(t,u) \in T \times U$, the set $M(t,u;\alpha)$ is nonempty and compact; and for all $(t,u) \in Z$, the set $M(t,u;\alpha)$ is convex,
- (A6) the correspondence $(t,u) \mapsto M(t,u;\alpha)$ is lower measurable,
- (A7) the correspondence $(t,u) \mapsto M(t,u;\alpha)$ is uniformly bounded by a p-integrable correspondence.[9]

That is, we assume standard conditions on the values of $M(\cdot;\alpha)$, along with a weak measurability property; moreover, we assume that choice sets do not become too large as we vary (t,u), in the sense that they are bounded by a p-integrable correspondence $\Upsilon: T \times U \rightrightarrows \mathfrak{R}^d$. Note, in particular, that we impose

[9]That is, there exists a lower measurable correspondence $\Upsilon: T \times U \rightrightarrows \mathfrak{R}^d$ with compact and convex values such that for all $\alpha \in \mathfrak{A}$ and all $(t,u) \in T \times U$, we have $M(t,u;\alpha) \subseteq \Upsilon(t,u)$, and $\int_{(t,u)} \sup \|\Upsilon(t,u)\|^p \mu(d(t,u)) < \infty$.

convexity on choice sets only on Z; if the probability measures $\{\nu(\cdot|t) : t \in T\}$ are non-atomic, then (A5) demands only that choice sets be nonempty and compact. Also note that if we strengthen (A3) to (A3′), assumption (A7) is implied by our other assumptions taking $\Upsilon(t,u) \equiv \text{co}A(t,u)$.

Finally, we impose the natural continuity assumption that choice sets are upper hemicontinuous on the space of p-integrable average choice functions. Henceforth, let $\mathfrak{A}^p = \{\alpha \in \mathfrak{A} : ||\alpha||_p < \infty\}$ denote the subset of p-integrable average choice functions, and endow $\mathfrak{A}^p$ with the topology inherited from the weak topology $\sigma(L^p, L^q)$ on $L^p \equiv L^p(T, \mathcal{T}, \kappa)$.[10] Assume:

(A8) for all $(t,u) \in T \times U$, the correspondence $\alpha \mapsto M(t,u;\alpha)$ has closed graph.

Thus, the degree p of integrability controls the tradeoff between our boundedness assumptions on $A(t,\cdot)$ and $M(\cdot;\alpha)$ (in (A3) and (A7)) and our continuity assumption on $M(t,u;\cdot)$ (in (A8)); of course, higher p strengthens boundedness and weakens continuity.

Main result. A choice function γ^* is a *choice equilibrium* if given the corresponding average choice function, γ^* is a selection from the corresponding choice sets; formally, letting $\alpha^* = \alpha(\cdot|\gamma^*)$, we require that $\gamma^*(t,u) \in M(t,u;\alpha^*)$ for each $(t,u) \in T \times U$. An *extremal choice equilibrium* is a choice equilibrium γ^* that selects from the extreme points of choice sets, i.e., for each $(t,u) \in T \times U$, we have $\gamma^*(t,u) \in \text{ext}M(t,u;\alpha^*)$, where $\text{ext}M(t,u;\alpha^*)$ is the set of extreme points of $M(t,u;\alpha^*)$.[11] We denote by $\overline{\text{ext}}M(t,u;\alpha^*)$ the closure of the set of extreme points of the choice correspondence. Our main theorem asserts existence of a choice equilibrium and a partial extremization result: given any choice equilibrium, there is a choice equilibrium such that choices are made from the closure of extreme points of choice sets for all (t,u) in the non-atomic part of ν and that is equivalent the choice equilibrium, in the sense that the

---

[10]Recall that $L^p(T,\mathcal{T},\kappa)$ is the set of κ-equivalence classes of $\mathcal{T}$-measurable mappings $\alpha: T \to \mathfrak{R}^d$ such that $||\alpha||_p \equiv \int_t ||\alpha(t)||^p \kappa(dt) < \infty$. Convergence in L^p is defined with reference to the dual space, $L^q \equiv L^q(T,\mathcal{T},\kappa)$, so that given any net $\{\alpha^\lambda\}$, we have $\alpha^\lambda \to \alpha$ if and only if for all $\varphi \in L^q$, we have $\int_t \varphi(t) \cdot [\alpha^\lambda(t) - \alpha(t)]\kappa(dt) \to 0$.

[11]For a (not necessarily convex) subset $X \subseteq \mathfrak{R}^d$, an element $x \in X$ is an *extreme point* of X if for all distinct $y, z \in X$, there does not exist $\beta \in (0,1)$ such that $x = \beta y + (1-\beta)z$.

equilibria determine the same average choices (and therefore same choice sets) up to a set of measure zero.

Theorem 2.1. *Assume (A1)–(A8). (a) A choice equilibrium exists; (b) for every choice equilibrium* γ^**, there exists a choice equilibrium* $\hat{\gamma}$ *such that (i)* γ^* *and* $\hat{\gamma}$ *determine equivalent average choices, i.e., for* κ*-almost all* t*,* $\alpha^*(t) = \hat{\alpha}(t)$*; and (ii)* $\hat{\gamma}$ *chooses from the closure of extreme points of choice sets for the non-atomic part of* ν*, i.e., for each* $(t,u) \in N \setminus Z$*, we have* $\hat{\gamma}(t,u) \in \overline{ext} M(t,u;\hat{\alpha})$*.*

Obviously, if the family $\{\nu(\cdot|t) : t \in T\}$ of conditional probabilities are non-atomic and the sets of extreme points are almost always closed, then extremal choice equilibria exist, and we can strengthen part (b) of the theorem to obtain a full purification result. Closedness of the set of extreme points does not hold generally (see Figure 7.4 of Aliprantis and Border (2006)), but it does hold widely; in particular, it holds if feasible alternative sets are simplicial, in which case the sets of extreme points are actually finite.

Corollary 2.2. *Assume that (A1)–(A8) hold; that* $Z = \varnothing$*; and that for each* $\alpha \in \mathfrak{A}$ *and for each* $(t,u) \in T \times U$*,* $ext M(t,u;\alpha)$ *is closed. (a) An extremal choice equilibrium exists; (b) for every choice equilibrium, there exists an extremal choice equilibrium* $\hat{\gamma}$ *that determines equivalent average choices, i.e., for* κ*-almost all* t*,* $\alpha^*(t) = \hat{\alpha}(t)$*.*

3. Large Games

In this section, we formulate a class of large games as a special case of the abstract framework. We endow the set N of players with a product structure, so that a player is described by a general characteristic t and a personal characteristic u, where the latter are distributed independently conditional on t in the space of players. Letting T and U be complete, separable metric spaces with Borel sigma-algebras $\mathcal{T}$ and $\mathcal{U}$, a *product large game* is described by a tuple $(T, U, A, P, \kappa, \nu)$ such that

- $N = T \times U$ is the player space,
- $A \colon N \rightrightarrows \mathfrak{R}^d$ is the feasible action correspondence,
- $P \colon N \times \mathfrak{M} \rightrightarrows \mathfrak{R}^d \times \mathfrak{R}^d$ is the preference correspondence,
- κ is a Borel probability measure on T,

- $\nu: T \times \mathcal{U} \to [0,1]$ is a transition probability,

where κ is the distribution of general characteristic t in the player space; $\mathfrak{M}$ is the set of κ-equivalence classes of measurable functions mapping T to $\mathfrak{R}^d$; and the transition probability $\nu(\cdot|t)$ gives the distribution of personal characteristics u conditional on t. Denote a generic player by $i = (t,u) \in N$. We set $\mathcal{N} = \mathcal{T} \otimes \mathcal{U}$, and we let $\mu = \nu(\cdot|\cdot) \otimes \kappa$ denote the joint distribution of players, as in the abstract framework. Setting Z as in Section 2, we make the following basic assumption to connect the large game model to Section 2:

(L1) assumptions (A1), (A2), and (A3′) from Section 2 hold.

A *strategy profile* is an $\mathcal{N}$-measurable mapping $\sigma: N \to \mathfrak{R}^d$ such that $\sigma(i) \in A(i)$ for all players i. Given a strategy profile σ, the implied average action is the function $\alpha: T \to \mathfrak{R}^d$ satisfying $\alpha(t) = \int_u \sigma(t,u)\nu(du|t)$ for κ-almost all t (where we identify functions equivalent up to sets of measure zero). Let $\mathfrak{A}$ denote the space of average actions. We interpret the set $P(i;\alpha)$ as the strict preference relation of player i given α.[12] A possible interpretation is that players are characterized by their personal characteristics, u, and by the groups to which they belong, t, and externalities (or "societal responses") are captured by the average actions α across groups; thus, the influence of society on a player's preferences, given a strategy σ, is captured by the infinite dimensional externality α.

Let $R: N \times \mathfrak{M} \rightrightarrows \mathfrak{R}^d \times \mathfrak{R}^d$ denote the weak preference correspondence corresponding to P, i.e., $R(i;\alpha) = \{(a',a) \in \mathfrak{R}^d \times \mathfrak{R}^d : (a,a') \notin P(i;\alpha)\}$. Fix player i and externality α. For each action a, let

$$\begin{aligned} P(i,a;\alpha) &= \{a' \in \mathfrak{R}^d : (a',a) \in P(i;\alpha)\} \\ R(i,a;\alpha) &= \{a' \in \mathfrak{R}^d : (a',a) \in R(i;\alpha)\} \\ R^{-1}(i,a;\alpha) &= \{a' \in \mathfrak{R}^d : (a,a') \in R(i;\alpha)\} \end{aligned}$$

denote the sets of actions that are strictly preferred to a, weakly preferred to a, and weakly worse than a for player i given α. We say $P(i;\alpha)$ is *irreflexive* if for all $a \in \mathfrak{R}^d$, we have $a \notin P(i,a;\alpha)$. Following Duggan (2011), we say that a set $Y \subseteq A(i)$ is *finitely dominant* if it is finite and for all $x \in A(i)$, there

[12] We remark that the formulation of the preference correspondence P is not subject to the critique in Balder (2000) (see Martins-da-Rocha and Topuzu (2008)).

exists $y \in Y$ with $y \in P(i,x;\alpha)$. Given $v \in A(i)$, we say that $P(i;\alpha)$ is *finitely subordinated* to v if there is a finitely dominant set Y with $v \in Y$ and such that there exists $z \in Y$ with $v \in P(i,z;\alpha)$ and $Y \backslash \{v\} \subseteq R^{-1}(i,z;\alpha)$. We say that $P(i;\alpha)$ satisfies the *finite-subordination property* if there is no $v \in A(i)$ such that $P(i;\alpha)$ is finitely subordinated to v.

We make the following further assumptions, which will later ensure that the players' best response sets are nonempty and satisfy weak measurability and continuity properties:

(L2) for all $i \in N$ and all $\alpha \in \mathfrak{A}$, $P(i;\alpha)$ is irreflexive and satisfies the finite-subordination property,
(L3) for all $i \in Z$, all $a \in A(i)$, and all $\alpha \in \mathfrak{A}$, $R(i,a;\alpha) \cap A(i)$ is convex,
(L4) for all $\alpha \in \mathfrak{A}$, the correspondence $i \mapsto R(i;\alpha)$ is lower measurable,
(L5) for all $i \in N$, the set $\{(a,\alpha) \in A(i) \times \mathfrak{A} : P(i,a;\alpha) \neq \varnothing\}$ is open in the product topology, where $\mathfrak{A}$ is endowed with the weak topology.

For all $i \in N$ and all $\alpha \in \mathfrak{A}$, define the correspondence $M \colon N \times \mathfrak{A} \rightrightarrows \mathfrak{R}^d$ by

$$M(i;\alpha) \quad = \quad \{a \in A(t,u) : P(t,u,a;\alpha) \cap A(i) = \varnothing\},$$

which consists of the maximal feasible actions for player i given externality α, i.e., it is player i's best response set. By Theorem 1 of Duggan (2011), the conditions (L1), (L2), and (L5) imply that $M(i;\alpha) \neq \varnothing$. And by a measurable maximum lemma (Duggan (2012b)), (L1) and (L4) imply that for all $\alpha \in \mathfrak{A}$, the correspondence $i \mapsto M(i;\alpha)$ is lower measurable.

A strategy profile σ^* is a *Nash equilibrium* if $\alpha^*(t) = \int_u \sigma^*(t,u)\nu(du|t)$ for κ-almost all t and $\sigma^*(i) \in M(i;\alpha^*)$ for each $i \in N$. This is readily seen as the specialization of a choice equilibrium from Section 2. The next result applies Theorem 2.1 to establish existence and extremization of Nash equilibrium. The proof consists in the straightforward verification that (L1)–(L5) imply (A1)–(A8).

Proposition 3.1. *Assume (L1)–(L5). (a) A Nash equilibrium exists; (b) for every Nash equilibrium σ^*, there exists a Nash equilibrium $\hat{\sigma}$ such that (i) σ^* and $\hat{\sigma}$ determine equivalent externalities, i.e., for κ-almost all t, $\alpha^*(t) = \hat{\alpha}(t)$; and (ii) $\hat{\sigma}$ chooses from the closure of extreme points of choice sets for the non-atomic part of ν, i.e., for each $i \in N \setminus Z$, we have $\hat{\sigma}(i) \in \overline{ext} M(i;\hat{\alpha})$.*

The first part of Proposition 3.1 does not explicitly require mixing and therefore provides a pure strategy existence result. We do need, however, some convexity for atomic players: for $i \in Z$, $a \in A(i)$, and $\alpha \in \mathfrak{A}$, (L3) implies $R(i,a;\alpha) \cap A(i)$ is convex. Such convexity obtains, e.g., if $A(i)$ is a simplex and each player's preferences are linear in her action, with $a \in A(i)$ representing a mixture over a finite set of pure actions. Viewed in this light, Proposition 3.1 may implicitly use mixed strategies. The second part of the proposition goes beyond this: when players' personal characteristics are atomless, we can set $Z = \varnothing$, so we do not require convexity of $A(i)$, and there is no implicit use of mixing. In fact, we find that every Nash equilibrium is equivalent to an extremal equilibrium, which allows us to make a stronger claim about existence of a pure strategy equilibrium in the simplicial model. Specifically, assuming personal characteristics are atomless and interpreting actions as mixtures over finite pure action sets, Proposition 3.1 implies that every Nash equilibrium is equivalent to a pure strategy equilibrium, and in particular, it establishes existence of equilibrium in pure strategies.

To formalize these ideas, we impose the following additional structure on the product large game model.[13]

- (L1′) (A2) holds,
- (L2′) for all $i \in N$, $A(i)$ is a non-empty face of the unit simplex in $\mathfrak{R}^d$,
- (L3′) for all $i \in N$ and all $\alpha \in \mathfrak{A}$, $P(i;\alpha)$ has a linear numerical representation $\pi_i(\cdot;\alpha)\colon \mathfrak{R}^d \to \mathfrak{R}$; in particular, for all $a, a' \in \mathfrak{R}^d$, we have $(a,a') \in P(i;\alpha)$ if and only if $\pi_i(a;\alpha) > \pi_i(a';\alpha)$.

The next corollary, an immediate implication of Proposition 3.1, formalizes the existence and purification results discussed above.

Corollary 3.2. *Assume (L1′)–(L3′), (L4), and (L5). Assume that for all $t \in T$, $\nu(\cdot|t)$ is atomless. (a) There exists a Nash equilibrium σ that chooses from the vertices of feasible action sets, i.e., for each $i \in N$, we have $\sigma(i) \in$*

[13] Note that we now assume feasible action sets are faces of the unit simplex, with elements representing mixtures over an underlying set of pure actions corresponding to the vertices of the simplex, and we assume linearity of payoffs with respect to mixtures over own actions. For purposes of existence of pure strategy equilibrium, we could change (L2′) so that feasible action sets $A(i)$ are subsets of vertices and drop (L3′), but we allow for mixing in order to state our result on purification.

ext$A(i)$; *(b) For every Nash equilibrium* σ^**, there exists a Nash equilibrium* $\hat{\sigma}$ *such that (i)* σ^* *and* $\hat{\sigma}$ *determine equivalent externalities, i.e., for* κ*-almost all* t*,* $\alpha^*(t) = \hat{\alpha}(t)$*; and (ii)* $\hat{\sigma}$ *chooses from the vertices of feasible action sets.*

Remark 3.3. Corollary 3.2 generalizes Theorem 2 and Remark 2 in Schmeidler (1973) from the model with a finite number of groups to the general model with a continuum of groups. More importantly, the assumption that best responses depend on the distribution across groups of average actions within groups, rather than the overall average action, puts Proposition 3.1 and Corollary 3.2 in an intermediate status compared to other results in the literature on large games with externalities captured by integrals. The externality α is an infinite-dimensional object, as opposed to the finite-dimensional externalities found in the literature, either the overall average action as in Rath (1992), or the finite-dimensional image of a function of the overall average action, as in Balder (2002), Martins-da-Rocha and Topuzu (2008), and Yu and Zhu (2005).[14] So we allow for players to respond to a much richer set of "societal variables," weakening considerably the implied notion of anonymity. Because we restrict the analysis to finite-dimensional action sets, whereas the literature allows for arbitrary compact action sets, our result is intermediate. □

Remark 3.4. Proposition 3.1 and Corollary 3.2 occupy an omitted position in Table 1 of Khan, Rath, and Sun (1997): the rightmost column of that table indicates that in games with uncountable action spaces and infinite-dimensional externalities, there is no pure-strategy Nash equilibrium. Here, we do have an uncountable action space and infinite-dimensional externalities, but the product structure of the player space, together with the result of Artstein (1989), allow us to work around the failure of Lyapunov's theorem in infinite dimensions, without having to move into saturated measure spaces with super-atomless measures.[15] □

[14]More precisely, the former two papers allow for infinite-dimensional externalities on the purely atomic part of the player space, while the externality on the non atomic part must be finite-dimensional. With non-atomicity, externalities do not have any infinite-dimensional component.

[15]See, e.g., Podczeck (2009), Carmona and Podczeck (2009,2014), and Khan, Rath, Sun, and Yu (2013). Carmona and Podczeck (2009) show that when externalities are the overall average, equilibria under super-nonatomicity are equivalent to equilibria with a Lebesgue

Remark 3.5. We capture finite-dimensional externalities analogous to those considered by Balder (2002) by restricting the influence of α on $P(i;\alpha)$. For simplicity we focus on the case in which κ is atomless. For each $j = 1, \ldots, m$, let $g_j: T \times \mathfrak{R}^d \to \mathfrak{R}$ be jointly measurable in (t,a) and continuous in a, and assume g_j is bounded by an integrable function of t. Then define the externality by

$$d(\alpha) = \left(\int_t g_j(t,\alpha(t))\kappa(dt) \right)_{j=1}^{m}.$$

Assume that $P(i;\alpha)$ has a numerical representation $\pi_i: \mathfrak{R}^d \times \mathfrak{R}^m \to \mathfrak{R}$ such that for all $a, a' \in \mathfrak{R}^d$, we have $(a,a') \in P(i;\alpha)$ if and only if $\pi_i(a; d(\alpha)) > \pi_i(a'; d(\alpha))$. For the degenerate case in which the personal characteristic u is fixed across all players, this reduces to the externality function of Balder (2002); and more generally, it extends his model to the case where the contribution of players with general characteristic t to the externality depends on the average actions of those players. □

Remark 3.6. Observe that Balder (2002) indexes players by t and assumes best response sets are convex for every atomic player, while our abstract framework requires convexity at every (t,u) pair such that u is an atom conditional on t; in this sense, our convexity assumption is more demanding than his. This difference can be traced back to our general form of externality. To see this, suppose that t is uniformly distributed on the unit interval $[0,1]$, and that the idiosyncratic variable is restricted to a single value, $\overline{u}$. Because the joint distribution on $(t,\overline{u})$ pairs is given by the uniform distribution on t (and is thus atomless), Balder's approach does not impose any convexity. In contrast, we do assume convexity, for in this setting integration over the idiosyncratic variable is trivial, and $\alpha(t) \equiv \sigma(t,\overline{u})$. This means that player i's preferences $P(i;\alpha)$ are general functions of the strategy profile itself, which is an infinite-dimensional object belonging to L^p. In lieu of the assumption of finite-dimensional externalities described in Remark 3.5, equilibrium existence thus relies on the stronger convexity assumption. □

Remark 3.7. We use (A3′) in Proposition 3.1, strengthening (A3), as a convenient way to obtain (A7). Under this stronger assumption, we can allow $p = \infty$, endow $\mathfrak{A}^\infty$ with the weak* topology, and directly obtain compactness of

---

space of players and a finite action space. Our approach allows us to have uncountable action spaces and a non-atomic (in particular Lebesgue) space of players.

the range of the correspondence S (defined formally in Appendix A) by noting that it belongs to a ball of finite radius in the L^∞-norm and use Alaoglu's theorem, instead of the indirect approach taken in Lemma A.7. □

Next, we return to the example of static competition among firms to illustrate some of the ideas above.

Example 3.8. *Cournot Game among Many Firms.* Recall the static firm competition model described in the Introduction. A firm $i = (t, u)$ is characterized by its location t and its technology u, and it produces a vector $q(i) \in Q(i) \subseteq \mathfrak{R}^d$. Profits $\pi_i(q;\alpha)$ depend on the production vector q and on the aggregate production vector at each location t, i.e., $\alpha(t) \equiv \int_u q(t,u)du$. Assume locations and technologies belong to complete, separable metric spaces, T and S; production sets $Q(i)$ are nonempty and compact, and the correspondence $i \mapsto Q(i)$ is lower measurable and integrably bounded; profits are jointly measurable in (i, q, α) and continuous in (q, α); and technologies u are nonatomically distributed. It is straightforward to verify that the conditions for our results are met,[16] so that a Nash equilibrium exists. To illustrate Corollary 3.2, consider the special case in which there are only $d > 0$ possible pure production plans for each firm, and feasibility correspondence $Q\colon N \rightrightarrows \mathfrak{R}^d$ is such that $Q(i)$ is a non-empty face of the unit simplex in $\mathfrak{R}^d$. Assume also that each firm i maximizes expected profit, so $\pi_i(\cdot;\alpha)$ is a linear function of mixtures over production plans for each aggregate production function α. Then Corollary 3.2 ensures the existence of a pure strategy equilibrium in the strong sense described above: for each firm i, $\sigma(i)$ selects exactly one of the d pure production plans. □

4. Stochastic Games

In this section, we consider the class of discounted stochastic games studied in Duggan (2012a) and map it into the framework of the current paper. Letting T and U be complete, separable metric spaces with Borel sigma-algebras $\mathcal{T}$

[16]Given that preferences are represented by the profit function, all that is left to be verified is that $i \mapsto M(i;\alpha)$ is lower measurable for each α and $\alpha \mapsto M(i;\alpha)$ is upper hemicontinuous for each i, where $M(i;\alpha) = \arg\max_{q\in Q(i)} \pi_i(q;\alpha)$. These follow from the maximum theorem and its measurable version (Aliprantis and Border (2006), Theorems 17.31 and 18.19).

and $\mathcal{U}$, a *noisy stochastic game* is a tuple $(T, U, \kappa, \nu_T, \nu_U, (X_i, B_i, \pi_i, \delta_i)_{i=1}^m)$ with $\{1, \dots, m\}$ being the set of players such that:

- $T \times U$, with sigma-algebra $\mathcal{T} \otimes \mathcal{U}$, is the measurable space of states,
- κ is a Borel probability measure on T,
- $\nu_T : T \times U \times X \times \mathcal{T} \to [0, 1]$ is a transition probability, with $X = \prod_i X_i$,
- $\nu_U : T \times \mathcal{U} \to [0, 1]$ is a transition probability,

and for each player $i = 1, \dots, m$,

- X_i is the action space,
- $B_i : T \times U \rightrightarrows X_i$ is the feasible action correspondence,
- $\pi_i : T \times U \times X \to \mathfrak{R}$ is the payoff function,
- $\delta_i \in [0, 1)$ is the discount factor.

These components describe a discrete-time dynamic game such that after any history, the relevant details of the game are described by a state (t, u), where t is a general component of the state that evolves as a stochastic function of the current state and actions, and u is a noise component. Specifically, timing is as follows: each period begins with a state (t, u), players simultaneously choose actions from their feasible sets $B_i(t, u)$; given action profile $x = (x_1, \dots, x_m)$, a new t' is realized from the distribution $\nu_T(\cdot|t, u, a)$, pinning down the general component of the next state; and the noise component u' of the new state is realized from $\nu_U(\cdot|t')$, which is independent of the state and actions in the previous period. Note that although u does not directly affect the transition on next period's noise component, it does enter into the transition on t, the feasible action correspondences B_i, and payoff functions π_i in a general way, and thus it is a non-trivial state variable.

We assume standard measurability properties imposed on stochastic games, including the assumption that the distribution of states is absolutely continuous with respect to a fixed measure, and that it is appropriately continuous in actions. Moreover, we assume that the distribution of the noise component is atomless:

(S1) $\nu_t(\cdot|t, u, x)$ is absolutely continuous with respect to κ for all $(t, u, x) \in T \times U \times X$,

(S2) for κ-almost all t, $\nu_U(\cdot|t)$ is absolutely continuous with respect to a fixed, atomless Borel probability measure λ on U, (S3) for all $(t,u) \in T \times U$, the mapping $x \to \nu_T(\cdot|t,u,x)$ is norm-continuous.

Let $g(\cdot|t,u,x)$ denote the density of $\nu_T(\cdot|t,u,x)$ with respect to κ, and let $h(\cdot|t)$ denote the density of $\nu_U(\cdot|t)$ with respect to λ. We impose the following weak integrability assumption on transition densities:

(S4) for all $(t,u,x) \in T \times U \times X$, the density $g(\cdot|t,u,x)$ belongs to $L^q(T,\mathcal{T},\kappa)$, i.e., $\int_{t'} |g(t'|t,u,x)|^q \kappa(dt') < \infty$, for some fixed $q > 1$.

Moreover, for all $i = 1,\dots,m$, we assume the standard compactness and metrizability of action sets, weak measurability of feasible action correspondences, and measurability and continuity properties of payoffs:

(S5) X_i is a compact metric space,
(S6) B_i is lower measurable with nonempty, compact values,
(S7) π_i is bounded and Carathéodory, i.e., $\pi_i(t,u,x)$ is measurable in (t,u) and continuous in x.

Letting $\Delta(X_i)$ denote the space of Borel probability measures on X_i with the weak* topology, a *stationary Markov strategy* for player i is a measurable mapping $\sigma_i : T \times U \to \Delta(X_i)$ such that for all $(t,u) \in T \times U$, $\sigma_i(t,u)$ assigns probability one to $B_i(t,u)$, i.e., $\sigma_i(B_i(t,u)|t,u) = 1$. Let σ denote a profile $(\sigma_i)_{i=1}^m$ as well as the transition probability $\otimes_{i=1}^m \sigma_i$ from states to action profiles. Continuation values $v_i(\cdot;\sigma)$ are uniquely defined by the recursion

$$\begin{aligned} v_i(t,u;\sigma) = (1-\delta_i) \int_x \pi_i(t,u,x)\sigma(dx|t,u) \\ + \delta_i \int_x \int_{t'} \int_{u'} v_i(t',u';\sigma)\nu_U(du'|t')\nu_T(dt'|t,u,x)\sigma(dx|t,u), \end{aligned}$$

and a *stationary Markov perfect equilibrium* is a profile σ such that for each $i = 1,\dots,m$ and each $(t,u) \in T \times U$, $\sigma_i(\cdot|t,u)$ places probability one on the solutions to

$$\max_{x_i \in B_i(t,u)} \int_{x_{-i}} \Bigg[(1-\delta_i)\pi_i(t,u,x) + \delta_i \int_{t'} \int_{u'} v_i(t',u';\sigma)\nu_U(du'|t')\nu_T(dt'|t,u,x) \Bigg] \sigma_{-i}(dx_{-i}|t,u).$$

Duggan (2012a) shows that under (S1)–(S3) and (S5)–(S7), a stationary Markov perfect equilibrium exists; we give an independent proof of existence adding (S4) below.

It will be convenient for us to define an auxiliary game, as follows. Given continuation values $v_i(\cdot;\sigma)$, let $\alpha_i(t) \equiv \int_u v_i(t,u;\sigma)\nu_U(du|t)$ denote the interim continuation value function generated by σ. For a profile $\alpha = (\alpha_i)_{i=1}^m$, we define the mapping $U(\cdot;\alpha): T \times U \times X \to \mathfrak{R}^m$ by

$$U_i(t,u,x;\alpha) \quad = \quad (1-\delta_i)\pi_i(t,u,x) + \delta_i \int_{t'} \alpha_i(t')g(t'|t,u,x)\kappa(dt')$$

for each $i = 1,\dots,m$. Let $\Gamma(t,u;\alpha)$ be the strategic form game with payoffs given by $U_i(\cdot;\alpha)$ for each player i. Let ξ denote both a profile $(\xi_i)_{i=1}^m$ of mixed strategies in $\Gamma(t,u;\alpha)$ and the corresponding product measure $\otimes_{i=1}^m \xi_i$, let $N(t,u;\alpha)$ denote the set of mixed strategy equilibrium profiles in $\Gamma(t,u;\alpha)$, and define

$$M(t,u;\alpha) \quad = \quad \left\{ \int U(t,u,x;\alpha)\xi(dx) : \xi \in N(t,u;\alpha) \right\},$$

the set of mixed strategy equilibrium payoff vectors of the auxiliary game $\Gamma(t,u;\alpha)$.

The next result applies Theorem 2.1 to establish existence and extremization of stationary Markov perfect equilibria. Note that the equilibrium payoffs in auxiliary games $\Gamma(t,u;\alpha)$ need not be convex, so we maintain the non-atomicity assumption (S2) throughout.[17]

Proposition 4.1. *Assume (S1)–(S7). (a) A stationary Markov perfect equilibrium exists; (b) for every stationary Markov perfect equilibrium σ^*, there exists another such equilibrium $\hat{\sigma}$ such that (i) σ^* and $\hat{\sigma}$ determine the same interim continuation values, i.e., for κ-almost all t and each $i = 1,\dots,m$, $\alpha_i^*(t) = \int_u v_i(t,u;\sigma^*)\nu_U(du|t) = \int_u v_i(t,u;\hat{\sigma})\nu_U(du|t) = \hat{\alpha}_i(t)$, and (ii) $\hat{\sigma}$ determines payoffs in the closure of extreme points of equilibrium payoffs in auxiliary games, i.e., for each $(t,u) \in T \times U$, we have $U(t,u,\hat{\sigma}(t,u);\hat{\alpha}) \in \overline{ext} M(t,u;\hat{\alpha})$.*

To apply Theorem 2.1 in the present setting, we transform the noisy stochastic game model into the abstract framework of Section 2 to obtain a choice equilibrium by part (a) of the theorem; we then transform the choice equilibrium back to the stochastic game model, and this delivers a stationary Markov perfect equilibrium. For the extremization result, we begin with a stationary

[17]Observe that because average actions are uniformly bounded, strengthening (A3′), we could allow $p = \infty$ and endow $\mathfrak{A}^\infty$ with the weak* topology. Then (S4) can be dropped, relying on the fact that the density $g(\cdot|t,u,x)$ is trivially integrable with respect to κ.

Markov perfect equilibrium in the stochastic game, then transform this into a choice equilibrium, and then apply part (b) of the theorem to deduce a payoff equivalent equilibrium that selects from the closure of extreme points of payoffs from auxiliary games. To elaborate on the first transformation, we set $d = m$, and we define the correspondence A of feasible alternatives as the correspondence $A(t,u) = [-\overline{\pi}, \overline{\pi}]^m$, where $\overline{\pi}$ bounds the absolute value of the players' payoff functions. An average action in the transformed model is a measurable mapping $\alpha: T \to \mathfrak{R}^m$ such that for each $t \in T$, $\alpha(t)$ represents a payoff vector averaged across the noise component u. We then define $U_i(t,u,x;\alpha)$ as above, and we let $M(t,u;\alpha)$ denote the set of mixed strategy equilibrium payoff vectors of the auxiliary game $\Gamma(t,u;\alpha)$. After verifying assumptions (A1)–(A8), Theorem 2.1 yields a choice equilibrium γ such that for all (t,u), $\gamma(t,u) \in M(t,u;\alpha)$, i.e., γ selects from mixed strategy equilibrium payoffs of auxiliary games, and from this (after addressing measurability issues), we back out a stationary Markov perfect equilibrium strategy profile of the original stochastic game.

The preceding result has sharp implications for noisy stochastic games such that in each state there is a single decision maker. A *noisy stochastic game with sequential moves* is a noisy stochastic game for which there exists a measurable mapping $p: T \times U \to \{1, \ldots, m\}$ such that for all $(t,u) \in T \times U$ and all $i \neq p(t,u)$, we have $|B_i(t,u)| = 1$. Thus, player $p(t,u)$ is the only player with a non-trivial choice of action in state (t,u), and we can therefore suppress the choices of other players notationally: let $B(t,u) = B_i(t,u)$ be the set of feasible actions of the active player $i = p(t,u)$ in state (t,u), and write payoffs $\pi_j(t,u,x_i)$, $j = 1, \ldots, m$, and transitions $\nu_T(\cdot|t,u,x_i)$ as functions of the actions x_i of the active player $i = p(t,u)$ only. The definition of equilibrium simplifies as well: then σ is a stationary Markov perfect equilibrium if for each $(t,u) \in T \times U$, the strategy $\sigma_i(\cdot|t,u)$ of the active player $i = p(t,u)$ places probability one on solutions to

$$\max_{x_i \in B(t,u)} (1-\delta_i)\pi_i(t,u,x_i) + \delta_i \int_{t'} \int_{u'} v_i(t',u';\sigma)\nu_U(du'|t')\nu_T(dt'|t,u,x_i).$$

Letting α be the interim continuation value generated by the stationary Markov strategy profile σ, the above maximization problem becomes

$$\max_{x_i \in B(t,u)} U_i(t,u,x_i;\alpha),$$

and we denote by $N^\bullet(t,u;\alpha)$ the set of solutions of this problem. Then the set $N(t,u;\alpha)$ consists of the mixtures over optimal actions for player $i = p(t,u)$ given σ, i.e., $N(t,u;\alpha) = \Delta(N^\bullet(t,u;\alpha))$, and $M(t,u;\alpha)$ is the set of payoff vectors generated by those mixtures over optimal actions.

The next proposition, which builds on Proposition 4.1, establishes existence of a stationary Markov perfect equilibrium in pure strategies and provides a purification result. Formally, a stationary Markovian strategy σ_i is *pure* if for all $(t,u) \in T \times U$, the probability measure $\sigma_i(t,u)$ puts probability one on a single action $x_i \in B_i(t,u)$. In the context of a game with sequential moves, we may view a pure strategy simply as a mapping from $T \times U$ to X_i and specify only the action of the active player; thus, a pure strategy profile is simply a measurable mapping $s: T \times U \to \bigcup_i X_i$ such that for all $(t,u) \in T \times U$, we have $s(t,u) \in B(t,u)$. Letting α be the interim continuation value generated by pure strategy profile s, we say s is a *pure stationary Markov perfect equilibrium* if for all $(t,u) \in T \times U$, we have $s(t,u) \in N^\bullet(t,u;\alpha)$.

Proposition 4.2. *Given a noisy stochastic game with sequential moves, assume (S1)–(S7). (a) A pure stationary Markov perfect equilibrium exists; (b) for every stationary Markov perfect equilibrium σ^*, there exists a pure stationary Markov perfect equilibrium $\hat{s}$ such that σ^* and $\hat{s}$ determine the same interim continuation values.*

Remark 4.3. The existence result of Nowak and Raghavan (1992) uses correlated strategies to obtain a form of stationary Markov perfect equilibrium, but in stochastic games with sequential moves, correlation can be removed from their equilibrium; thus, their theorem implies existence of stationary Markov perfect equilibria in stochastic games with sequential moves. Proposition 4.2 shows that in the noisy version of their model, we can go one step further and obtain equilibria in pure strategies. □

The next example illustrates the results of this section in an infinite-horizon Stackelberg model of oligopolistic competition.

Example 4.4. *Sequential Oligopoly Game with Random Movers.* Consider a discrete time dynamic oligopoly model, with infinite horizon and discounting. Instead of a large number of firms as in Example 3.8, there is a finite number $m > 1$ of firms. Existence of stationary Markov perfect equilibria in mixed strategies for the model with simultaneous moves follows from Proposition

4.1. Here, however, we consider a version of the model in which firms move sequentially and establish existence of equilibrium in pure strategies. The general component t of a state determines the firm that is to move at that state, and it also determines other stochastic properties of the industry (for instance, the firms' capital stock levels); conditional on t, the noise component u is independently and non-atomically distributed, determining, e.g., firm-specific technological shocks, aggregate demand shocks, and/or aggregate labor supply shocks. Let X_i be the compact metric space of possible production plans for each firm $i = 1, \dots, m$. Assume that the game satisfies (S1)–(S6). Observe that ν_t will necessarily have atoms as t determines the active firm at the state. Let B_i be the feasibility correspondence, which is a singleton for the non-active firms at state (t, u), and lower measurable, with non-empty and compact values. Let $\pi_i(t, u, x)$ be the Carathéodory payoff function of firm i, which depends only on the action of the active firm at a state, fulfilling (S7), so Proposition 4.2 ensures existence of a pure stationary Markov perfect equilibrium. Drawing a parallel with the second part of Example 3.8, we allow for a general space of pure production plans X_i (rather than a finite set), obtain a stationary equilibrium where the actions of the active firm are possibly mixed, and then use our extremization results to conclude that there is an equivalent equilibrium where the active firm chooses pure production plans only. □

5. Endogenous Institutions

In this section, we consider a general framework for endogenous selection of individuals into groups (e.g., institutions, jurisdictions, or clubs) and collective decision within groups. We seek equilibria such that each of a continuum of individuals selects into her optimal group, given the distribution of other individuals across groups and collective decisions within the groups, and such that the collective decision of each group is consistent with the membership of that group. Letting T and U be complete, separable metric spaces with Borel sigma-algebras $\mathcal{T}$ and $\mathcal{U}$, a *product society with endogenous institutions* is given by a tuple $(T, U, G, C, Y, \Pi, \kappa, \nu, (P_j)_{j=1}^n)$ such that

- $N = T \times U$ is the space of individuals,
- $G = \{1, \dots, n\}$ is a set of possible groups,

- $C: N \rightrightarrows G$ is a membership constraint correspondence,
- Y is the set of possible group decisions,
- $\Pi: N \times G \times Y^n \times \mathfrak{B} \to \mathfrak{R}$ is the utility function of individual i,
- for each $j = 1, \ldots, n$, $P_j: Y \times \mathfrak{A} \rightrightarrows Y$ is the group decision correspondence,
- κ is a Borel probability measure on T,
- $\nu: T \times \mathcal{U} \to [0,1]$ is a transition probability,

where $\mathfrak{A}$ is the space of κ-equivalence classes of $\mathcal{T}$-measurable mappings $\alpha: T \to \mathfrak{R}^n$ such that for κ-almost all t, we have $\alpha_j(t) \geq 0$ for all $j = 1, \ldots, n$ and $\sum_{j=1}^n \alpha_j(t) = 1$, and $\mathfrak{B}$ is the space of κ-equivalence classes of $\mathcal{T}$-measurable mappings $\beta: T \to \mathfrak{R}$ such that for κ-almost all t, $\beta(t) \in [0,1]$. We let $\mu = \nu(\cdot|\cdot) \otimes \kappa$ denote the joint distribution over individuals.

We interpret the elements of this model as follows: an individual $i = (t,u)$ is described by a public characteristic t and private characteristic u; $C(i)$ is the set of groups that individual i may join; $\beta(t)$ represents the fraction of individuals with public characteristic t who join a given group, so that the payoff $\Pi(i,j,y,\beta)$ reflects i's preferences over the group j to which she belongs, the vector $y = (y_1, \ldots, y_n)$ of group decisions, and the membership β within her group; the function α summarizes the distribution of public characteristics within each group; the correspondence P_j determines a set $P_j(y,\alpha)$ of viable group decisions for group j as a function of group decisions and group membership;[18] and κ gives the distribution of public characteristics within society, and $\nu(\cdot|t)$ is the conditional distribution of private characteristics.

We impose weak assumptions of nonemptiness and measurability of the membership constraint correspondence and non-atomicity of the distribution of private characteristics:

(I1) for all $i \in N$, $C(i)$ is nonempty,

(I2) the correspondence $C: N \rightrightarrows G$ has measurable sections, i.e., for each group j, the set $\{i \in N : j \in C(i)\}$ is $\mathcal{T} \otimes \mathcal{U}$-measurable,

(I3) Y is a compact subset of $\mathfrak{R}^m$,

[18] Caplin and Nalebuff (1997) are somewhat ambiguous regarding the domain of their group decision functions; in their footnote 23, they mention that P_j is defined for vectors $(\tau_1, \ldots, \tau_n)$ that contain in each component a measure τ_j with $\tau_j(T) \leq 1$ describing the selection of public characteristics into each group. Our formulation is more useful for present purposes.

(I4) for κ-almost all t, $\nu(\cdot|t)$ is absolutely continuous with respect to a fixed, atomless Borel probability measure λ on U.

We endow $\mathfrak{B}$ with the topology inherited from the weak topology of κ-integrable equivalence classes of $\mathcal{T}$-measurable mappings $\beta: T \to \mathfrak{R}$. Moreover, assume utility functions possess standard measurability and continuity properties:

(I5) Π is Carathéodory, i.e., $\Pi(i, j, y, \beta)$ is $\mathcal{T} \otimes \mathcal{U}$-measurable in $i = (t, u)$ and continuous in (j, y, β),

and for each $j = 1, \ldots, n$, assume the following convexity and continuity properties of group decisions:

(I6) P_j has non-empty, convex values,
(I7) P_j has closed graph.

A *strategy profile* is a $\mathcal{T} \otimes \mathcal{U}$-measurable mapping $s: N \to G$ such that $s(i) \in C(i)$ for all $i \in N$, where we interpret $s(i)$ as the group that individual i joins. A *membership profile* is a mapping $\alpha \in \mathfrak{A}$, where $\alpha(t) = (\alpha_1(t), \ldots, \alpha_n(t))$ gives the fraction $\alpha_j(t)$ of individuals with public characteristic t who join each group j. Then an *endogenous institutional equilibrium* is a triple (s, α, y) such that

(i) for all i, $s(i)$ solves

$$\max_{j \in C(i)} \Pi(i, j, y, \alpha_j)$$

(ii) for all $j = 1, \ldots, n$ and κ-almost all t,

$$\alpha_j(t) = \nu(\{(t, u) \in T \times U : s(t, u) = j\}|t),$$

(iii) for all $j = 1, \ldots, n$, $y_j \in P_j(y, \alpha)$.

In words, the individuals choose optimally, given the vector y of group decisions and membership profile α; the membership profile reflects the actual selection of public characteristics into the groups; and the decision of each group is consistent with the group membership and decisions of other groups.

The next result, which follows from Theorem 2.1, establishes the existence of endogenous institutional equilibria and generalizes the analysis of Section 6 of Caplin and Nalebuff (1997).

Proposition 5.1. *Assume (I1)–(I7). There exists an endogenous institutional equilibrium.*

The proof proceeds by transforming the product society model into our abstract framework, adding an artificial agent for each group, and applying our abstract results to deduce a choice equilibrium; we then transform the choice equilibrium back to the original model to obtain an endogenous institutional equilibrium. We specify the expanded set $\tilde{T} = T \cup G$ of public characteristics, adding one element for each group and topologizing $\tilde{T}$ in the obvious way,[19] and we select an arbitrary $\tilde{u} \in U$ to stand for the private characteristic of the artificial agents. We set $d = n + m$, where the first n coordinates are used to track the selection of individuals into groups, and the last m coordinates are used to capture the decisions of artificial agents. We associate each group j with the unit coordinate vector e^j in $\mathfrak{R}^{n+m}$, and we define $A(t,u) = \text{co}\{e^j : j \in C(i)\}$ for each individual $i = (t,u) \in T \times U$. For an artificial agent j, we set $A(j,\tilde{u}) = \{0_n\} \times Y$, where 0_n denotes a vector of n zeroes. An average action in the transformed model is a measurable mapping $\tilde{\alpha}: \tilde{T} \to \mathfrak{R}^{n+m}$ such that for each $t \in T$, $\tilde{\alpha}(t)$ gives the fraction of individuals with public characteristic t who join group j, and for each $j \in G$, $\tilde{\alpha}(j)$ specifies a decision for group j. To define choice sets, for $(t,u) \in T \times U$, we specify that $M(t,u;\tilde{\alpha})$ consists of the convex hull of unit coordinate vectors e^j in $\mathfrak{R}^{n+m}$ determined by solutions to

$$\max_{j \in C(i)} \Pi(t,u,j,\tilde{\alpha}(1),\ldots,\tilde{\alpha}(n),\tilde{\alpha}_j|_T),$$

and for $j \in G$, we specify

$$M(j,\tilde{u};\tilde{\alpha}) \quad = \quad \{0_n\} \times P_j(\tilde{\alpha}(1),\ldots,\tilde{\alpha}(n),\tilde{\alpha}_1|_T,\ldots,\tilde{\alpha}_n|_T).$$

In particular, given a selection of individuals into groups and the decisions of other groups, we require that each artificial agent j essentially chooses an action from the set of decisions determined by the decision correspondence for group j. Then assumptions (A1)–(A8) are satisfied, and Theorem 2.1 yields a choice equilibrium $\hat{\gamma}$ such that for all $(t,u) \in T \times U$, $\hat{\gamma}(t,u)$ is an extreme point of $A(t,u)$, and this mapping to extreme points yields an equilibrium strategy profile s, as required.

[19]We give $\tilde{T}$ a metric that is equal to the original metric restricted to T and such that each singleton $\{j\}$ is open.

Remark 5.2. We have formulated our equilibrium concept in terms of pure strategies, following the literature. We can extend our equilibrium concept to allow for mixing, as in the product large game model of Section 3. To do so, we would reformulate the model so that each group j is associated with the unit coordinate vector in $\Re^n$, and we would define a mixed strategy profile $\sigma: N \to \Delta_n$ from individuals to the unit simplex in $\Re^n$, where $\sigma(i)$ represents the mixture over groups used by individual i. We then define (σ, α, y) as a *mixed endogenous institutional equilibrium* by replacing (i) and (ii) in the original definition by

(i′) for all i, $\sigma(i)$ puts probability one on solutions to

$$\max_{j \in C(i)} \Pi(i, j, y, \alpha_j),$$

(ii′) for all $j = 1, \ldots, n$ and κ-almost all t,

$$\alpha_j(t) \quad = \quad \int_u \sigma_j(t) \nu(du|t).$$

With this definition, we can apply our extremization result in Theorem 2.1 to conclude that every mixed endogenous institutional equilibrium, σ, is equivalent to a pure one, s, in the sense that σ and s determine the same membership profile, and therefore the same group decisions. □

Remark 5.3. An issue considered in the literature on local public goods with mobility is the existence of an equilibrium in which non-trivial sorting into groups occurs, where there are multiple non-empty groups with distinct distributions of types among their members. See Gomberg (2004) for further discussion of this issue. Because we allow an individual to have direct preferences over groups and the membership within her group, the endogenous institutional equilibria of our model will typically possess this feature, but in a simplified version of the model with utilities defined only over group decision vectors, e.g., $\Pi(i, y)$, it may be that in equilibria with multiple nonempty groups, all nonempty groups are identical. Thus, we do not address the problem of non-trivial sorting in such simplified settings. □

An interesting application of the model is that in which group decisions are determined by voting. We assume that given individual selection into groups, the voting rule in group j is given by a quota $q_j \in [\frac{1}{2}, 1)$, and we assume

(I8) for κ-almost all t, all y, all α, and all distinct $j,k \in G$,

$$\nu\left(\{u \in U : \Pi(t,u,j,y,\alpha_j) = \Pi(t,u,k,y,\alpha_k)\} \middle| t\right) \;=\; 0.$$

That is, given public characteristic t, group decisions y, and group membership α, the set of individuals who are indifferent between any two groups is measure zero. A special case is when private characteristics $u \in \mathfrak{R}^n$ consist of additive, continuously distributed shocks to the utility from group membership, as in Caplin and Nalebuff (1997). We posit a correspondence $F_j\colon Y \times \mathfrak{A} \rightrightarrows Y$ for each group j such that $F_j(y,\alpha)$ consists of the feasible decisions for the group, and we assume:

(I9) for each $j = 1,\ldots,n$, F_j has non-empty, closed values and is continuous.

Given public characteristic t, group decision vector y, and membership profile α, let

$$U(t,j,y,\alpha) \;=\; \left\{u \in U : j \in \arg\max_{k \in C(t,u)} \Pi(t,u,k,y,\alpha_k)\right\}$$

be the set of private characteristics such that it is optimal for individual (t,u) to join group j. Furthermore, given $z_j, z_j' \in Y$, let

$$V(z_j', z_j | t,j,y,\alpha) \;=\; \left\{u \in U(t,j,y,\alpha) : \begin{array}{c} \Pi(t,u,j,(z_j',y_{-j}),\alpha_j) \\ > \Pi(t,u,j,(z_j,y_{-j}),\alpha_j) \end{array}\right\}$$

be the subset of private characteristics u such that (t,u) joins group j and strictly prefers z_j' to z_j. We then define the dominance relation $D(j,y,\alpha)$ on Y such that $z_j' D(j,y,\alpha) z_j$ holds if and only if

$$\int_t \nu(V(z_j', z_j|t,j,y,\alpha)|t)\kappa(t) \;>\; q_j \int_t \alpha_j(t)\kappa(dt),$$

i.e., a fraction of members of group j greater than the quota q_j strictly prefers z_j' to z_j. Finally, we define $P_j(y,\alpha)$ as the set of feasible decisions that are maximal for this dominance relation; formally, we assume:

(I10) for all j, all y, and all α,

$$P_j(y,\alpha) \;=\; \{z_j \in F_j(y,\alpha) : \forall z_j' \in F_j(y,\alpha), \neg z_j' D(j,y,\alpha) z_j\}.$$

We refer to $((T,\mathcal{T}),(U,\mathcal{U}),G,C,Y,\Pi,\kappa,\nu,(F_j,q_j)_{j=1}^n)$ as a *product society with endogenous voting institutions.*

The next lemma establishes the closed graph property in (I7), which is key for existence in the voting setting.

Lemma 5.4 *Given a product society with endogenous voting institutions, conditions (I1)–(I5), and (I8)–(I10) together imply (I7).*

It is well-known that existence of voting equilibria can be problematic, depending on the dimensionality of the set Y of group decisions relative to the quota q_j. More can be said, however, if we impose additional weak structure on voting games.

(I11) for each $j = 1, \dots, n$, F_j has convex values,
(I12) $\Pi(i, j, y, \alpha_j)$ is quasi-concave in y_j.

It is known that under (I9)–(I12), if $m < \frac{1}{1-q_j}$, then the set $P_j(y, \alpha)$ of group decisions is nonempty, as required by (I6). Convexity of these sets can be problematic in multiple dimensions, but in one dimension, since $m = 1$, we of course have $m < \frac{1}{1-q_j}$; and moreover, the sets $P_j(y, \alpha)$ are thus non-empty, compact intervals. Thus, conditions (I1)–(I7) are satisfied, and an equilibrium exists.

We end with an example that illustrates the scope of Proposition 5.1 in the context of a one-dimensional model of local public goods with mobility that allows for externalities in consumption and production of public good and for crowding effects within jurisdictions.[20]

Example 5.5. *Super Majority Voting over Local Public Goods.* Assume that a continuum of voters freely select into n fixed jurisdictions, that $t \in \mathfrak{R}_+$ indicates an individual's wealth and u is a preference parameter, and that within each jurisdiction j a one-dimensional public good level y_j is determined and financed by a proportional tax on wealth of its residents. Let $c_j(y)$ be the cost of producing y_j units of public good when other jurisdictions produce $y_{-j} = (y_1, \dots, y_{j-1}, y_{j+1}, \dots, y_n)$, and we assume that c_j is continuous, that it is convex and strictly increasing in y_j, and that $y_j = 0$ implies $c_j(y) = 0$. Let $r(\alpha_j) = \int_t \alpha_j(t)\kappa(dt) + \epsilon$ denote the tax base of jurisdiction j given membership profile α, where for simplicity we assume a positive baseline $\epsilon > 0$ of revenues.

[20]Existence of equilibrium in multidimensional spaces under voting by quota rule is considered in Greenberg and Shitovitz (1988) without mobility and by Konishi (1996) with mobility.

The set of public good levels feasible for jurisdiction j is then

$$F_j(y,\alpha) \quad = \quad \{z_j \in \mathfrak{R}_+ : c_j(z_j, y_{-j}) \leq r(\alpha_j)\},$$

which has non-empty, compact, convex values and is continuous. Given public good levels $y = (y_1, \ldots, y_n)$ and tax rate $\theta_j \in [0,1]$, we assume the utility of voter $i = (t,u)$ is quasilinear in disposable wealth, measurable, and given by $\pi(j,y,u,\alpha_j) + (1-\theta_j)t$, where π is increasing and concave in y_j and continuous in (y,α_j). Taxes are determined by a budget balance constraint, so we can write an individual's utility from locating in jurisdiction j as

$$\Pi(t,u,j,y,\alpha_j) \quad = \quad \pi(j,y,u,\alpha_j) + \left(1 - \frac{c_j(y)}{r(\alpha_j)}\right)t,$$

which is continuous in (y,α_j) and quasi-concave in y_j, and we impose the additional restriction that for all $t \in T$, all $j \in G$, all $\alpha_j \in \mathfrak{B}$, all distinct $y, y' \in \mathfrak{R}^n_+$ and all $d \in \mathfrak{R}$, we have $\nu(\{u : \pi(j,y,u,\alpha_j) - \pi(j,y',u,\alpha_j) = d\}|t) = 0$, so that indifference between jurisdictions holds among a measure zero set of voters. We assume that voting in each district j is subject to a quota $q_j \in [\frac{1}{2}, 1)$, which may be majority rule, $q_j = \frac{1}{2}$. Conditions (I1)–(I5), and (I8)–(I10) are satisfied, and by Lemma 5.4, so is (I7). Since $m = 1$ and voter preferences over public good levels in their district are continuous and single-peaked, fulfilling (I11) and (I12), we have (I6) as well. Therefore, Proposition 5.1 implies the existence of an endogenous institutional equilibrium, in which each voter locates optimally given public good levels and taxes within jurisdictions, and the public good levels and taxes are consistent with the quota rule used in each jurisdiction. □

Appendix A. Proof of Main Result

We assume throughout this section of the appendix that conditions (A1)–(A8) hold. We make use of the following measurability property of the norm of the feasible set of alternatives. Henceforth, we abbreviate references to Aliprantis and Border (2006) as AB.

Lemma A.1. *The mapping* $(t,u) \to \sup \|A(t,u)\|$ *is* $\mathcal{N}$*-measurable.*

Proof. Note three observations: with continuity of the Euclidean norm $\|\cdot\|$, (A2) implies that the correspondence $(t,u) \mapsto \|A(t,u)\|$ is lower measurable with nonempty, closed values in $\mathfrak{R}$; as a consequence, there is a sequence

$\{f_n\}$ of $\mathcal{N}$-measurable functions $f_n: N \to \mathfrak{R}$ such that for all (t, u), $\|A(t, u)\| = \mathrm{cl}\{f_n(t, u)\}$ (AB, Corollary 18.14); and the pointwise limit of a sequence of measurable functions into a complete, separable metric space is itself measurable (AB, Lemma 4.29). Therefore, $(t, u) \mapsto \sup \|A(t, u)\| = \sup\{f_n(t, u)\}$ is $\mathcal{N}$-measurable. □

The next lemma establishes existence of average choice functions that are p-integrable; the existence of such functions does not follow from (A3), because that assumption does not restrict feasible action sets across t, but it does follow from (A4)–(A7).

Lemma A.2. *$\mathfrak{A}^p$ is nonempty.*

Proof. Since A is lower measurable with closed values, it admits a measurable selection γ (Theorem 18.13 in AB), and then $\alpha = \alpha(\cdot|\gamma)$ is an average choice function, and $\alpha \in \mathfrak{A}$. Then (A5) and (A6) imply that $M(\cdot; \alpha)$ admits a measurable selection $\tilde{\gamma}$; (A4) implies that $\tilde{\alpha} = \alpha(\cdot|\tilde{\gamma})$ is an average choice function; and (A7) implies $\|\tilde{\alpha}\|_p < \infty$. Therefore, $\tilde{\alpha} \in \mathfrak{A}^p$. □

Proceeding to the main result, let us first define two useful correspondences. Let $A^*: T \rightrightarrows \mathfrak{R}^d$ be defined by

$$A^*(t) = \int_u A(t, u)\nu(du|t),$$

and for each $\alpha \in \mathfrak{A}$, define $M^*(\cdot; \alpha): T \rightrightarrows \mathfrak{R}^d$ by

$$M^*(t; \alpha) = \int_u M(t, u; \alpha)\nu(du|t).$$

These are, respectively, the Aumann integrals of the feasible action correspondence $A(t, \cdot)$, and of the choice correspondence $M(t, \cdot; \alpha)$, with respect to u. The next lemma characterizes the average choice functions in terms of the correspondence A^*, and it characterizes the almost everywhere selections from $M(\cdot; \beta)$ (for any given average choice function β) in terms of $M^*(\cdot; \beta)$.

Lemma A.3. *For each $\mathcal{T}$-measurable $\alpha: T \to \mathfrak{R}^d$, (a) α is a κ-almost everywhere selection from A^* if and only if $\alpha \in \mathfrak{A}$; (b) for each $\beta \in \mathfrak{A}^p$, α is a κ-almost everywhere selection from $M^*(\cdot; \beta)$ if and only if there exists a choice function γ with $\gamma(t, u) \in M(t, u; \beta)$ for μ-almost all (t, u) such that for κ-almost all t, $\alpha(t) = \int_u \gamma(t, u)\nu(du|t)$.*

Proof. To prove (a), note that the "if" direction is immediate from the definition of average choice function. Indeed, letting $\alpha \in \mathfrak{A}$ be determined as $\alpha = \alpha(\cdot|\gamma)$ for the choice function γ, it follows that for all $t \in T$, $\gamma(t,\cdot): U \to \mathfrak{R}^d$ is a selection from $A(t,\cdot): U \rightrightarrows \mathfrak{R}^d$, and therefore $t \mapsto \int_u \gamma(t,u)\nu(du|t)$ is a $\mathcal{T}$-measurable selection from A^*. Since $\alpha(t) = \int_u \gamma(t,u)\nu(du|t)$ for κ-almost all t, this direction is proved. For the "only if" direction, let α be a κ-almost everywhere selection from A^*. Then the theorem of Artstein (1989) yields a $\mathcal{T} \otimes \mathcal{U}$-measurable mapping $\gamma: T \times U \to \mathfrak{R}^d$ such that for κ-almost all t, we have: $\alpha(t) = \int_u \gamma(t,u)\nu(du|t)$, and for $\nu(\cdot|t)$-almost all u, $\gamma(t,u) \in A(t,u)$. In particular, his assumptions (i)–(vi) are fulfilled, respectively, by the assumptions that T and U are Polish, the assumption that $\nu(\cdot|\cdot): \mathcal{U} \times T \to [0,1]$ is a transition probability, (A1), (A2), and (A3). Thus, α is determined as $\alpha = \alpha(\cdot|\gamma)$ for the choice function γ. The proof of (b) is parallel, using $M(\cdot;\beta)$ and $M^*(\cdot;\beta)$ instead of A and A^*. □

Next, we verify pointwise properties of the correspondences A^* and $M^*(\cdot;\alpha)$.

Lemma A.4. *(a) the correspondence A^* has nonempty, compact, and convex values; (b) for each $\alpha \in \mathfrak{A}$, the correspondence $M^*(\cdot;\alpha)$ has nonempty, compact, and convex values.*

Proof. Nonemptiness in (a) follows from (A1) and (A2), which imply that A is lower measurable with nonempty, closed values, and so it admits a measurable selection (see Theorem 18.13 of AB), and from (A3), which implies that this selection is $\nu(\cdot|t)$-integrable for κ-almost all t. Nonemptiness in (b) follows similarly from (A3)–(A6). For compactness in (a) and (b), note that (A3) implies that for κ-almost all t, the correspondence $u \mapsto A(t,u)$ is p-integrably bounded with respect to $\nu(\cdot|t)$. By a version of Fatou's lemma (see Proposition 7 (p.73) of Hildenbrand (1974)), the integral $A^*(t) = \int_u A(t,u)\nu(du|t)$ of this correspondence is compact. Similarly, the integral $M^*(t;\alpha) = \int_u M(t,u;\alpha)\nu(du|t)$ is compact. Let $Z(t) = \{u \in U : (t,u) \in Z\}$, and note that for each $t \in T$,

$$A^*(t) \;=\; \int_{u \in U \smallsetminus Z(t)} A(t,u)\nu(du|t) + \int_{u \in Z(t)} A(t,u)\nu(du|t),$$

so convexity in (a) follows because the first term on the right-hand side is convex by non-atomicity of $\nu(\cdot|t)$ on $U \smallsetminus Z(t)$ and a version of Lyapunov's theorem (see Theorem 8.6.3 of Aubin and Frankowska (1990)), and because the second term is convex by (A1). The argument for convexity in (b) is parallel, using (A5) instead of (A1). □

The next lemma elaborates on convexity of $M^*(t;\alpha)$. Note that because the latter set is nested between the integral of extreme points of $M(t,u;\alpha)$ and the integral of the convex hull, the lemma in fact implies equality of all three sets. A further implication, since $M(t,u;\alpha)$ is convex for all $(t,u) \in Z$ by (A5), is that for each $\alpha \in \mathfrak{A}$ and each $t \in T$, we have $M^*(t;\alpha) = \int_u \text{co}M(t,u;\alpha)\nu(du|t)$. As in the preceding proof, we define $Z(t) = \{u \in U : (t,u) \in Z\}$ for all $t \in T$.

Lemma A.5. *For each $\alpha \in \mathfrak{A}$ and each $t \in T$, we have*

$$\int_{u \in U \setminus Z(t)} extM(t,u;\alpha)\nu(du|t) \quad = \quad \int_{u \in U \setminus Z(t)} coM(t,u;\alpha)\nu(du|t).$$

Proof. Fix $\alpha \in \mathfrak{A}$ and $t \in T$. Note that $u \mapsto M(t,u;\alpha)$ is p-integrably bounded by (A3) and (A4). Then

$$\begin{aligned}\int_{u \in U \setminus Z(t)} \text{ext}M(t,u;\alpha)\nu(du|t) &= \int_{u \in U \setminus Z(t)} \text{coext}M(t,u;\alpha)\nu(du|t) \\ &= \int_{u \in U \setminus Z(t)} \text{co}M(t,u;\alpha)\nu(du|t),\end{aligned}$$

where the first equality follows from non-atomicity of $\nu(\cdot|t)$ on $U \setminus Z(t)$ and a version of Lyapunov's theorem (see Theorem 8.6.3 of Aubin and Frankowska (1990)), and the second follows from (A5) and the observation that the convex hull of a compact subset of finite-dimensional Euclidean space is equal to the convex hull of its extreme points ($\text{co}C = \text{coextco}C$ by the Krein-Milman theorem, Theorem 7.68 of AB), and the fact that C and its convex hull possess the same extreme points ($\text{extco}C = \text{ext}C$). □

We require the domain of our fixed point argument to be convex and norm-closed.

Lemma A.6. *The set $\mathfrak{A}^p$ is convex and norm-closed.*

Proof. Convexity follows from Lemmas A.3 and A.4, which establish that $\mathfrak{A}$ consists of all measurable selections from the convex-valued correspondence A^*. To prove norm-closedness, assume the sequence $\{\alpha^m\}$ in $\mathfrak{A}^p$ converges to α in L^p, and therefore there is a subsequence (still indexed by m) and a κ-measure zero set $R \in \mathcal{T}$ such that for all $t \notin R$, $\alpha^m(t) \to \alpha(t)$ (AB, Theorem 13.6).

Given any $t \notin R$, since $\alpha^m(t) \in A^*(t)$ for all m, and since $A^*(t)$ is compact by Lemma A.4, it follows that $\alpha(t) \in A^*(t)$. Then Lemma A.3 yields $\alpha \in \mathfrak{A}^p$. $\square$

Define $S : \mathfrak{A}^p \rightrightarrows L^p$ so that $S(\alpha)$ consists of all $\mathcal{T}$-measurable, κ-almost everywhere selections from $M^*(\cdot;\alpha)$. The correspondence S will be the focus of our fixed point argument.

Lemma A.7. *The range of S, $S(\mathfrak{A}^p)$, is a relatively compact subset of $\mathfrak{A}^p$.*

Proof. Let S^* consist of all $\mathcal{T} \otimes \mathcal{U}$-measurable, μ-almost everywhere selections from Υ given in (A7). By Lemma A.4, S^* is compact in $L^p(T \times U, \mathcal{T} \times \mathcal{U}, \mu)$ endowed with the weak topology induced by $L^q(T \times U, \mathcal{T} \times \mathcal{U}, \mu)$, by a general version of Diestel's theorem (c.f. Balder (1990), Corollary 4.2) and Remark 3.1 in Yannelis (1991). Now define the mapping $\phi : S^* \to L^p$ by $\phi(\beta)(t) = \int_u \beta(t,u)\nu(du|t)$ for all $\beta \in S^*$. Indeed, the range is well-specified because

$$\int_t ||\phi(\beta)(t)||^p \kappa(dt) \le \int_{(t,u)} ||\beta(t,u)||^p \mu(d(t,u)) < \infty,$$

where the first inequality follows from Jensen's inequality and the second from (A7). We claim that ϕ is continuous. Consider a net $\{\beta^\nu\}$ in S^* such that $\beta^\nu \to \beta \in S^*$ in the weak topology, and consider any $\varphi \in L^q$. Then

$$\begin{aligned} &\int_t [\varphi(t) \cdot (\phi(\beta^\nu)(t) - \phi(\beta)(t))]\kappa(dt) \\ &= \int_{(t,u)} [\varphi(t) \cdot (\beta^\nu(t,u) - \beta(t,u))]\mu(d(t,u)) \\ &\to 0, \end{aligned}$$

as claimed. Therefore, $\phi(S^*)$ is a compact subset of L^p, and because $S(\alpha) \subseteq \phi(S^*)$ for all $\alpha \in \mathfrak{A}^p$, it follows that the range $S(\mathfrak{A}^p)$ is relatively compact. That it is a subset of $\mathfrak{A}^p$ follows from Lemma A.3 and the observation that $M^*(t;\alpha) \subseteq A^*(t)$ for all t. $\square$

Next, we verify pointwise properties of the correspondence S.

Lemma A.8. *S has nonempty, convex, and closed values.*

Proof. For nonemptiness, consider any $\alpha \in \mathfrak{A}^p$. By (A6), $M(\cdot;\alpha)$ admits a measurable selection γ (AB, Theorem 18.13). Defining the average choice function β by $\beta(t) = \int_u \gamma(t,u)\nu(du|t)$, Lemma A.3 implies $\beta \in S(\alpha)$. Convex values follow from Lemma A.4. To prove that $S(\alpha)$ is weakly closed in L^p, we use the previous observation that it is convex: then, by Theorem 13 (p.422) of Dunford and Schwartz (1957), it follows that $S(\alpha)$ is weakly closed if and only if it is norm closed. To show that it is norm closed, let $\{\beta^m\}$ be a sequence in $S(\alpha)$ that converges in the L^p-norm to β. Then there is a subsequence (still indexed by m) that converges pointwise κ-almost everywhere (AB, Theorem 13.6). Let $R^0 \subseteq T$ be a κ-measure zero Borel set such that $\beta^m(t) \to \beta(t)$ for all $t \in T \setminus R^0$, and for each m, let $R^m \subseteq T$ be a κ-measure zero Borel set such that $\beta^m(t) \in M^*(t;\alpha)$ for all $t \in T \setminus R^m$. Set $R = \bigcup_{k=0}^{\infty} R^k$, and note that $\kappa(R) = 0$. Given any $t \in T \setminus R$, we have $\beta^m(t) \in M^*(t;\alpha)$ for all m and $\beta^m(t) \to \beta(t)$. Since $M^*(\cdot;\alpha)$ has closed values, by Lemma A.4, it follows that $\beta(t) \in M^*(t;\alpha)$, and therefore $\beta \in S(\alpha)$, as required. □

The next lemma establishes that S is sequentially upper hemicontinuous, in the sense that for all weakly closed sets $F \subseteq L^p$, the lower inverse $S^\ell(F) = \{\alpha \in \mathfrak{A}^p : S(\alpha) \cap F \neq \varnothing\}$ is sequentially closed. Because the range of S is relatively compact in the weak topology by Lemma A.7, the Eberlein-Šmulian theorem (AB, Theorem 6.34) implies that relative compactness is equivalent to sequential relative compactness: every sequence in $S(\mathfrak{A}^p)$ has a weakly convergent sequence with limit in L^p. In light of this fact, sequential upper hemicontinuity reduces to sequentially closed graph.

Lemma A.9. *S is sequentially upper hemicontinuous.*

Proof. By Lemma A.7, S has relatively compact range, so we must show its graph is sequentially closed. Let $\{\alpha^m\}$ be a sequence converging to α in $\mathfrak{A}^p$, and let $\beta^m \in S(\alpha^m)$ for each m with $\beta^m \to \beta$. For each m, by Lemma A.3, there is a choice function γ^m such that for κ-almost every t, we have $\beta^m(t) = \int_u \gamma^m(t,u)\nu(du|t)$ and for μ-almost every (t,u), we have $\gamma^m(t,u) \in M(t,u;\alpha^m) \subseteq \Upsilon(t,u)$. Thus, there is a Borel set $Q^m \subseteq T \times U$ with $\mu(Q^m) = 0$ such that for all $(t,u) \in (T \times U) \setminus Q^m$, we have $\gamma^m(t,u) \in M(t,u;\alpha^m)$. Define $Q = \bigcup_{m=1}^{\infty} Q^m$, and note that $\mu(Q) = 0$. From the proof of Lemma A.7, we know that the set S^* of selections from Υ is compact. It follows that the set $\{\gamma^m\}$ is relatively compact. By the Eberlein-Šmulian theorem (AB, Theorem 6.34), the sequence has a subsequence (still indexed by m) that converges weakly to some $\gamma \in L^p(T \times U, \mathcal{T} \otimes \mathcal{U}, \mu)$. This subsequence $\{\gamma^m\}$ is uniformly integrable, and Proposition C of Artstein (1979) implies that for μ-almost all (t,u), the action $\gamma(t,u)$ is a convex combination of accumulation points of $\{\gamma^m(t,u)\}$. In particular, there is a Borel set $Q^* \supseteq Q$ with $\mu(Q^*) = 0$ such that for all $(t,u) \in (T \times U) \smallsetminus Q^*$, this property holds. Then we can let T^* be a Borel set with $\kappa(T^*) = 0$ such that for all $t \in T \smallsetminus T^*$, we have $\nu(\{u \in U : (t,u) \in Q^*\}|t) = 0$. Consider any $t \in T \smallsetminus T^*$. Then for $\nu(\cdot|t)$-almost all u, we have $(t,u) \notin Q^*$, and we can write $\gamma(t,u) \in \text{co}\{a^1,\ldots,a^k\}$, where for each a^j, $j = 1,\ldots,k$, there is a subsequence $\{\gamma^{m_\ell}\}$ such that $\gamma^{m_\ell}(t,u) \to a^j$. (The dependence of the subsequence on j is notationally suppressed.) Note that $\gamma^{m_\ell}(t,u) \in M(t,u;\alpha^{m_\ell})$ for each ℓ. By (A8), we have $a^j \in M(t,u;\alpha)$ for $j = 1,\ldots,k$, so $\gamma(t,u) \in \text{co}M(t,u;\alpha)$. By Lemma A.5, the $\mathcal{T}$-measurable mapping $\beta^*: T \to \mathfrak{R}^d$ defined by $\beta^*(t) = \int_u \gamma(t,u)\nu(du|t)$ is a κ-almost everywhere selection from $M^*(\cdot;\alpha)$, i.e., $\beta^* \in S(\alpha)$. Lastly, we claim that $\{\beta^m\}$ converges weakly to β^*. Indeed, consider any $f \in L^q$, and note that

$$\begin{aligned}
\int_t f(t)\beta^m(t)\kappa(dt) &= \int_t f(t)\left(\int_u \gamma^m(t,u)\nu(du|t)\right)\kappa(dt) \\
&= \int_{(t,u)} f(t)\gamma^m(t,u)\mu(d(t,u)) \\
&\to \int_{(t,u)} f(t)\gamma(t,u)\mu(d(t,u)) \\
&= \int_t f(t)\left(\int_u \gamma(t,u)\nu(du|t)\right)\kappa(dt) \\
&= \int_t f(t)\beta^*(t)\kappa(dt),
\end{aligned}$$

where the second and third equalities follow from Proposition 2.3.2 of Rao (2005), and the limit from weak convergence $\gamma^m \to \gamma$. Therefore, $\beta^m \to \beta^*$ weakly, and since the weak topology on L^p is Hausdorff, we conclude that $\beta = \beta^*$, which implies $\beta \in S(\alpha)$, as required. □

We can now complete the proof of the theorem. For (a), we observe that the correspondence $S: \mathfrak{A}^p \rightrightarrows L^p$ satisfies the conditions of Theorem 2.3 of Agarwal and O'Regan (2002). In particular, Lemma A.6 shows that $\mathfrak{A}^p$ is a nonempty, convex, norm-closed subset of L^p; Lemma A.8 shows that $S(\alpha)$ is nonempty, weakly closed, and convex for each $\alpha \in \mathfrak{A}^p$; Lemma A.9 shows that the correspondence is sequentially upper hemicontinuous; furthermore, Lemma A.7 shows that the range of S is a relatively compact subset of $\mathfrak{A}^p$ with the weak topology. Then there exists $\alpha^* \in \mathfrak{A}^p$ satisfying $\alpha^* \in S(\alpha^*)$. Since α^* is a selection from $M^*(\cdot;\alpha^*)$, Lemma A.3 yields an almost everywhere selection γ^1 from $M(\cdot;\alpha^*)$ such that α^* is determined by the choice function γ^1. To modify γ^1 on the μ-measure zero exceptional set $E = \{(t,u) \in T \times U : \gamma^1(t,u) \notin M(t,u;\alpha^*)\}$, note that since $M(\cdot;\alpha^*)$ is lower measurable with nonempty, closed values, we can take a measurable selection γ^2 of $M(\cdot;\alpha^*)$ by the Kuratowski-Ryll-Nardzewski theorem (AB, Theorem 18.13). Define the choice function γ^* by

$$\gamma^*(t,u) = \begin{cases} \gamma^1(t,u) & \text{if } (t,u) \in (T \times U) \setminus E \\ \gamma^2(t,u) & \text{if } (t,u) \in E. \end{cases}$$

This preserves average choices, as $\alpha^*(t) = \int_u \gamma^*(t,u)\nu(du|t) = \alpha(t|\gamma^*)$ for κ-almost all t, and γ^* is a choice equilibrium.

For part (b), let γ^* be a choice equilibrium with corresponding average choice function α^*. Define the correspondence $\hat{M}: T \times U \rightrightarrows \mathfrak{R}^d$ by

$$\hat{M}(t,u) = \begin{cases} \overline{\text{ext}} M(t,u;\alpha^*) & \text{if } (t,u) \notin Z \\ M(t,u;\alpha^*) & \text{else.} \end{cases}$$

By (A5) and Lemma C.1, it follows that $\overline{\text{ext}} M(\cdot;\alpha^*)$ is lower measurable with nonempty, compact values. As the measurable splicing of two such correspondences, $\hat{M}$ is lower measurable with nonempty, compact values. Since (A5) implies $\text{ext} M(t,u;\alpha^*) \subseteq \overline{\text{ext}} M(t,u;\alpha^*) \subseteq \text{co} M(t,u;\alpha^*)$, Lemma A.5 implies that for all $t \in T$,

$$\begin{aligned} &\int_u \hat{M}(t,u)\nu(du|t) \\ &= \int_{u \in U \setminus Z(t)} \overline{\text{ext}} M(t,u;\alpha^*)\nu(du|t) + \int_{u \in Z(t)} M(t,u;\alpha^*)\nu(du|t) \\ &= M^*(t;\alpha^*), \end{aligned}$$

and therefore for κ-almost all t, we have $\alpha^*(t) \in \int_u \hat{M}(t,u)\nu(du|t)$. The correspondence $\hat{M}$ satisfies the conditions of Artstein's (1989) theorem, and therefore there exists a $\mathcal{T} \otimes \mathcal{U}$-measurable mapping $\hat{\gamma}^1: T \times U \to \mathfrak{R}^d$ such that for κ-almost all t, we have: $\alpha^*(t) = \int_u \hat{\gamma}^1(t,u)\nu(du|t) = \alpha(t|\hat{\gamma}^1)$, and for $\nu(\cdot|t)$-almost all u, $\hat{\gamma}^1(t,u) \in \hat{M}(t,u)$. To modify $\hat{\gamma}^1$ on the μ-measure zero exceptional set $\hat{E} = \{(t,u) \in T \times U : \hat{\gamma}^1(t,u) \notin \hat{M}(t,u)\}$, note that since $\hat{M}$ is lower measurable with nonempty, closed values, we can take a measurable selection $\hat{\gamma}^2$ of $\hat{M}$ by the Kuratowski-Ryll-Nardzewski theorem (AB, Theorem 18.13). Define the choice function $\hat{\gamma}$ by

$$\hat{\gamma}(t,u) = \begin{cases} \hat{\gamma}^1(t,u) & \text{if } (t,u) \in (T \times U) \setminus \hat{E} \\ \hat{\gamma}^2(t,u) & \text{if } (t,u) \in \hat{E}. \end{cases}$$

This preserves average choices, as $\alpha^*(t) = \int_u \hat{\gamma}(t,u)\nu(du|t) = \alpha(t|\hat{\gamma})$ for κ-almost all t, and it provides the choice equilibrium called for by the theorem.

Appendix B. Proofs for Applications

This section of the appendix contains proofs of results for applications to large games, stochastic games, and endogenous institutions.

B.1. **Large Games.**

Proof of Proposition 3.1. Parts (a) and (b) follow from the correspondence between product large games and the abstract framework upon verifying (A1)–(A8). We have noted that assumptions (A1)–(A3′) follow directly from (L1), with (A3′) giving us (A7) as well. The definition of $M(i;\alpha)$ immediately yields (A4). Nonemptiness of $M(i;\alpha)$ follows from Theorem 1 of Duggan (2011); and compactness of $M(i;\alpha)$ follows immediately from the open graph assumption (L5). Given $i \in Z$, irreflexivity (from (L2)) and convexity (from (L3)) of $R(i,\cdot;\alpha)$ yield convexity of $M(i;\alpha)$. Indeed, suppose $M(i;\alpha)$ is not convex, so there exist distinct $x, y \in M(i;\alpha)$ and $\lambda \in (0,1)$ such that $z = \lambda x + (1-\lambda)y \notin M(i;\alpha)$. Then there exists $w \in A(i)$ such that $w \in P(i,z;\alpha)$. Since $x \in M(i;\alpha)$, we have $x \in R(i,w;\alpha)$, and similarly $y \in R(i,w;\alpha)$. But then convexity of $R(i,w;\alpha)$ implies $z \in R(i,w;\alpha)$, a contradiction. We conclude that (A5) holds. As mentioned above, the result in Duggan (2012b) implies (A6). Finally, (L5) implies (A8) by standard continuity arguments. □

B.2. **Stochastic Games.**

Proof of Proposition 4.1. Consider a noisy stochastic game, and set $N = T \times U$, $\mu = \nu_U(\cdot|\cdot) \otimes \kappa$, and $d = m$. Let $\overline{\pi}$ bound the absolute value of the players' payoff functions, and define the correspondence A of feasible alternatives by $A(t,u) = [-\overline{\pi}, \overline{\pi}]^m$. Given a choice function $\gamma: T \times U \to \mathfrak{R}^m$, we define the corresponding average choice function $\alpha: T \to \mathfrak{R}^m$ by $\int_u \gamma(t,u)\nu_U(du|t)$, and for technical reasons we place average action profiles in $\mathfrak{A}^p$, where p is conjugate to q from (S4), with the topology inherited from the weak topology on $L^p(T, \mathcal{T}, \kappa)$. Define

$$U_i(t,u,x;\alpha) \quad = \quad (1-\delta_i)\pi_i(t,u,x) + \delta_i \int_{t'} \alpha_i(t')g(t'|t,u,x)\kappa(dt'),$$

and let $M(t,u;\alpha)$ be the set of mixed strategy equilibrium payoff vectors of the auxiliary game $\Gamma(t,u;\alpha)$. Then (A1)–(A4) and (A7) are immediately implied by our specifications of A and M. Assumption (A5) is directly implied by continuity of $U(t,u,x;\alpha)$ in x and (S2), and (A6) follows from Lemma 2 of Duggan (2012a). For (A8), we use the argument in the proof of Lemma 1 of Duggan (2012a), with a slight modification (in addition to changes in notation): for a fixed (t,u), we take a sequence (x^n,α^n) in the graph of $M(t,u;\cdot)$ with $(x^n,\alpha^n)\to(x,\alpha)$ in $X\times\mathfrak{A}^p$, and we invoke (S4) to obtain

$$\int_{t'}(\alpha_i^n(t')-\alpha_i(t'))g(t'|t,u,x)\kappa(dt') \quad\to\quad 0,$$

and the remainder of the argument proceeds as in Lemma 1 of that paper, making use of (S3).

To prove part (a), we use part (a) of Theorem 2.1 to deduce existence of a choice equilibrium γ. Letting α denote the associated average choice function, it follows that γ is a selection from $M(\cdot;\alpha)$. Extending $U(t,u,\cdot;\alpha)$ to mixed profiles in the obvious way, Lemmas 1 and 2 of Duggan (2012a) imply that $U(t,u,\xi;\alpha)$ is Carathéodory in (t,u,ξ) and that the correspondence $(t,u)\mapsto N(t,u;\alpha)$ is lower measurable. Furthermore, for each $(t,u)\in T\times U$, there exists $\xi\in N(t,u;\alpha)$ such that $\gamma(t,u)=U(t,u,\xi;\alpha)$. Then by Filippov's implicit function theorem (AB, Theorem 18.17), there is a measurable mapping $\sigma\colon T\times U\to\prod_i\Delta(X_i)$ such that for all (t,u), we have $\sigma(t,u)\in N(t,u;\alpha)$ and $\gamma(t,u)=U(t,u,\sigma(t,u);\alpha)$. Then σ determines the same payoffs as γ and selects from mixed strategy equilibria of auxiliary games, and thus σ is a stationary Markov perfect equilibrium, as required.

For part (b), consider a stationary Markov perfect equilibrium σ, and define $\gamma\colon T\times U\to\mathfrak{R}^m$ by $\gamma(t,u)=(v_1(t,u;\sigma),\ldots,v_n(t,u;\sigma))$. The corresponding average choice function, defined by $\alpha_i(t)=\int_u v_i(t,u;\sigma)\nu_U(du|t)$, is readily seem as the interim continuation value function generated by σ, defined above. The best response problem of player i in state (t,u) simplifies to

$$\max_{x_i\in B_i(t,u)}\int_{x_{-i}}U_i(t,u,x;\alpha)\sigma_{-i}(dx_{-i}|t,u),$$

and therefore we have $\gamma(t,u)\in M(t,u;\alpha)$. Thus, the stationary Markov perfect equilibrium σ determines a choice equilibrium γ. Then part (b) of Theorem 2.1, with (S2), yields a choice equilibrium $\hat{\gamma}$ that determines the same average choice function and chooses from the closure of extreme points of $M(t,u;\alpha)$

for all (t,u). By the argument above, we can then use Filippov's implicit function theorem to construct a stationary Markov perfect equilibrium $\hat{\sigma}$ that determines the same payoffs as $\hat{\gamma}$, and thus it induces the same interim continuation values as σ and determines payoffs in the closure of extreme points of auxiliary games, as required. □

Proof of Proposition 4.2. Let σ^* be a stationary Markov perfect equilibrium, the existence of which follows from Proposition 4.1. By Proposition 4.1, there is a stationary Markov perfect equilibrium $\hat{\sigma}$ such that conditions (i) and (ii) of the proposition hold; in particular, $\hat{\sigma}$ determines payoffs in the closure of extreme points of the equilibrium payoffs in auxiliary games for every state. Let $\alpha^* = \hat{\alpha}$ be the corresponding interim continuation values. For $(t,u) \in T \times U$, let $i = p(t,u)$, and note that for every payoff vector $y \in \text{ext}M(t,u;\hat{\alpha})$, there exists $x_i \in N^\bullet(t,u;\hat{\alpha})$ such that $y = U(t,u,x_i;\hat{\alpha})$. Indeed, let $\xi_i \in N(t,u;\hat{\alpha})$ satisfy $y = U(t,u,\xi_i;\hat{\alpha})$. By Carathéodory's convexity theorem (AB, Theorem 5.32), we can choose this mixed action so that the support of ξ_i consists of $k \leq m+1$ pure actions, say $\{x_i^1,\ldots,x_i^k\}$.[21] Suppose $y \neq y^1 = U(t,u,x_i^1;\hat{\alpha})$. Since $y = U(t,u,\xi_i;\hat{\alpha})$, there is a mixture ξ_i' over $\{x_i^2,\ldots,x_i^k\}$ such that $y \neq y^2 = U(t,u,\xi_i';\hat{\alpha})$ and $y = \xi_i(x_i^1)y^1 + (1-\xi_i(x_i^1))y^2$. But $x_i^1, \xi_i' \in N(t,u;\hat{\alpha})$, so $y^1, y^2 \in M(t,u;\hat{\alpha})$, and $y^1 \neq y^2$, and then y is not an extreme point of $M(t,u;\hat{\alpha})$, a contradiction.

This observation extends more generally: we claim that for every payoff vector $y \in \overline{\text{ext}}M(t,u;\hat{\alpha})$, there exists $x_i \in N^\bullet(t,u;\hat{\alpha})$ such that $y = U(t,u,x_i;\hat{\alpha})$. To see this, let $\{y^n\}$ be a sequence in $\text{ext}M(t,u;\hat{\alpha})$ such that $y^n \to y$, and for each n, let $x_i^n \in N^\bullet(t,u;\hat{\alpha})$ satisfy $y^n = U(t,u,x_i^n;\hat{\alpha})$. Since $B(t,u)$ is compact, we can go to a subsequence (still indexed by n) such that $x_i^n \to x_i \in B(t,u)$. By continuity of $U(t,u,\cdot;\hat{\alpha})$ and the maximum theorem (AB, Theorem 17.31), we then have $y = U(t,u,x_i;\hat{\alpha})$ and $x_i \in N^\bullet(t,u;\hat{\alpha})$, as claimed. Lower measurability of the correspondence $(t,u) \mapsto N^\bullet(t,u;\hat{\alpha})$ can be shown by a simplification of the proof of Lemma 2 in Duggan (2012a). Then, using the facts that $U(t,u,x_i;\hat{\alpha})$ is Carathéodory in (t,u,x_i) and that for each $(t,u) \in T \times U$, there exists $x_i \in N^\bullet(t,u;\hat{\alpha})$ with $U(t,u,\sigma^*(t,u);\hat{\alpha}) = U(t,u,x_i;\hat{\alpha})$, Filippov's

[21]In fact, y belongs to the convex hull of $\{U(t,u,x_i;\hat{\alpha}) : x_i \in \text{supp } \xi_i\}$, so Carathéodory's theorem allows us to write y as a convex combination of $\{U(t,u,x_i^1;\hat{\alpha}),\ldots,U(t,u,x_i^k;\hat{\alpha})\}$. As player i is the only active player at state (t,u), we can replace ξ_i with the corresponding weights on $\{x_i^1,\ldots,x_i^k\}$.

implicit function theorem (AB, Theorem 18.17) yields a measurable mapping $\hat{s}: T \times U \to \bigcup_i X_i$ such that $\hat{s}(t,u) \in N^{\bullet}(t,u;\hat{\alpha})$ and $U(t,u,\sigma^*(t,u);\hat{\alpha}) = U(t,u,\hat{s}(t,u);\hat{\alpha})$, for all (t,u). Then $\hat{s}$ is a pure stationary Markov perfect equilibrium that determines the same continuation values as $\hat{\sigma}$, and therefore the same interim continuation values as σ^*. □

B.3. **Endogenous Institutions.**

Proof of Proposition 5.1. We specify the set $\tilde{T} = T \cup G$ of general characteristics, adding one element for each group and topologizing $\tilde{T}$ in the obvious way. Specifically, if ρ denotes the original metric on T, we define $\tilde{\rho}$ on $\tilde{T}$ so that $\tilde{\rho}|_T = \rho$, and for all distinct $\tilde{t}, t' \in \tilde{T}$, if $\tilde{t} \in G$, then $\tilde{\rho}(\tilde{t}, t') = 1$, and we extend κ to a Borel probability measure $\tilde{\kappa}$ on $\tilde{T}$ by letting κ' be uniform on G and setting $\tilde{\kappa} = \frac{1}{2}\kappa + \frac{1}{2}\kappa'$. We fix an arbitrary $\tilde{u} \in U$, and we extend the transition probability ν to $\tilde{\nu}: \tilde{T} \times \mathcal{U} \to [0,1]$ such that for all $t \in T$, $\tilde{\nu}(\cdot|t) = \nu(\cdot|t)$, and for all $j \in G$, $\tilde{\nu}(\cdot|j)$ is degenerate on $\tilde{u}$. Setting $d = n + m$, we associate each group j with the unit coordinate vector e^j in $\mathfrak{R}^{n+m}$, and we define $A(t,u) = \mathrm{co}\{e^j : j \in C(i)\}$ for each individual $i = (t,u) \in T \times U$. For an artificial agent j, we set $A(j,\tilde{u}) = \{0_n\} \times Y$, where 0_n denotes a vector of n zeroes. Thus, (A1)–(A3′) are satisfied.

An average action in the transformed model is a measurable mapping $\tilde{\alpha}: \tilde{T} \to \mathfrak{R}^{n+m}$, where for each $t \in T$, $\tilde{\alpha}(t)$ is an $(n+m)$-tuple such that the first n coordinates give the fraction of individuals with general characteristic t who join group j, with zeroes in the last m coordinates; and for each $j \in G$, $\tilde{\alpha}(j)$ has zeroes in the first n coordinates and is such that the last m coordinates comprise a vector in Y.

To define choice sets, note that the projection of the restriction of $\tilde{\alpha}$ onto the first n coordinates, i.e., $(\tilde{\alpha}_1|_T, \ldots, \tilde{\alpha}_n|_T)$, is a membership profile for individuals in the original model. For $(t,u) \in T \times U$, we then specify that

$$M(t,u;\tilde{\alpha}) \quad = \quad \mathrm{co}\left\{ e^j : j \in \arg\max_{j \in C(i)} \Pi(t,u,j,\tilde{\alpha}(1),\ldots,\tilde{\alpha}(n),\tilde{\alpha}_j|_T) \right\},$$

and for $j \in G$, we specify

$$M(j,\tilde{u};\tilde{\alpha}) \quad = \quad \{0_n\} \times P_j(\tilde{\alpha}(1),\ldots,\tilde{\alpha}(n),\tilde{\alpha}_1|_T,\ldots,\tilde{\alpha}_n|_T).$$

In particular, given a selection of individuals into groups and the decisions of other groups, we require that each artificial agent j essentially chooses an action from the set of decisions determined by the decision correspondence for group j. Then (A4), (A5), and (A7) are immediately satisfied. To obtain (A6), note that the objective function $\Pi(t, u, j, \tilde{\alpha}(1), \ldots, \tilde{\alpha}(n), \tilde{\alpha}_j|_T)$ is Carathéodory, i.e., measurable in (t, u) and continuous in j, so a measurable version of the maximum theorem (AB, Theorem 18.19) implies that the correspondence of maximizers is lower measurable. Then $M(\cdot; \tilde{\alpha})$, as the pointwise convex hull of a lower measurable correspondence, is itself lower measurable. Finally, (A8) follows from continuity of $\Pi(i, j, y, \alpha_j)$ in (y, α_j) and the theorem of the maximum.

Thus, Theorem 2.1 yields a choice equilibrium $\hat{\gamma}$ such that for all $(t, u) \in T \times U$, $\hat{\gamma}(t, u)$ is an extreme point of $A(t, u)$. Let $\hat{\alpha}$ be the average action determined by $\hat{\gamma}$, and let $\alpha = (\hat{\alpha}_1|_T, \ldots, \hat{\alpha}_n|_T)$ be the projection onto the first n coordinates of the restriction of $\hat{\alpha}$ to T. As a choice equilibrium, it follows that for all $j \in G$, we have $\hat{\gamma}(j, \tilde{u}) \in M(j, \tilde{u}; \hat{\alpha})$; or equivalently, setting $y_j = \hat{\gamma}(j, \tilde{u})$ for each group j, we have $y_j \in P_j(y, \alpha)$. It also follows that for all $(t, u) \in T \times U$, $\hat{\gamma}(t, u)$ is a vertex of $A(t, u)$ corresponding to a solution of the problem,

$$\max_{j \in C(i)} \Pi(i, j, y, \alpha_j),$$

of individual $i = (t, u)$. Finally, we define the strategy profile s as follows: for all $i = (t, u) \in T \times U$, we identify the group j such that $\hat{\gamma}(t, u) = e^j$, and we specify $s(i) = j$. The triple (s, α, y) so-defined satisfies conditions (i)–(iii), and we conclude that there exists an endogenous institutional equilibrium, as required. □

Proof of Lemma 5.4. Let $\{(y^k, \alpha^k, z_j^k)\}$ be a sequence converging to (y, α, z_j) in $Y \times \mathfrak{A} \times Y$ such that for all k, $z_j^k \in P_j(y^k, \alpha^k)$. Suppose that $z_j \notin P_j(y, \alpha)$. Since F_j is upper hemi-continuous, we have $z_j \in F(y, \alpha)$, and so there exists $z'_j \in F_j(y, \alpha)$ such that $z'_j D(j, y, \alpha) z_j$. Then we have

$$\int_t \nu(V(z'_j, z_j | t, j, y, \alpha) | t) \kappa(t) \quad > \quad q_j \int_t \alpha_j(t) \kappa(dt).$$

Since F_j is lower hemi-continuous, there is a sequence $\{\tilde{z}_j^k\}$ such that $\tilde{z}_j^k \in F_j(y^k, \alpha^k)$ for all k and $\tilde{z}^k \to z'_j$. Of course, we have $\int_t \alpha_j^k(t) \kappa(dt) \to \int_t \alpha_j(t) \kappa(dt)$. By (I8), note that for κ-almost all t, $\nu(\cdot|t)$-almost all $u \in U(t, j, y, \alpha)$, and all $\ell \neq j$, we have

$$\Pi(t, u, j, y, \alpha_j) \quad > \quad \Pi(t, u, \ell, y, \alpha_\ell).$$

Thus, for μ-almost all (t,u) with $u \in U(t,j,y,\alpha)$, continuity of $\Pi(t,u,\cdot)$ implies

$$\Pi(t,u,j,y^k,\alpha_j^k) \quad > \quad \Pi(t,u,\ell,y^k,\alpha_\ell^k)$$

for sufficiently large k, and in particular, $u \in U(t,j,y^k,\alpha^k)$ for high k. Furthermore, continuity of $\Pi(t,u,\cdot)$ implies that for μ-almost all (t,u) such that $u \in V(z_j', z_j|t,j,y,\alpha)$, we have $u \in V(\tilde{z}_j^k, z_j^k|t,j,y^k,\alpha^k)$ for high k. Let $S^k = \{(t,u) : u \in V(\tilde{z}_j^k, z_j^k|t,j,y^k,\alpha^k)\}$ and $S = \{(t,u) : u \in V(z_j', z_j|t,j,y,\alpha)\}$. By Fatou's lemma, we then have

$$\begin{aligned}\liminf_{k\to\infty} \int_t \nu(V(\tilde{z}_j^k, z_j^k|t,j,y^k,\alpha^k)|t)\kappa(dt) &= \liminf_{k\to\infty} \int_{(t,u)} I_{S^k}(t,u)\mu(d(t,u)) \\ &\geq \int_{(t,u)} \liminf_{k\to\infty} I_{S^k}(t,u)\mu(d(t,u)) \\ &\geq \int_{(t,u)} I_S(t,u)\mu(d(t,u)) \\ &= \int_t \nu(V(z_j', z_j|t,j,y,\alpha)|t)\kappa(t).\end{aligned}$$

Thus, we have

$$\int_t \nu(V(\tilde{z}_j^k, z_j^k|t,j,y^k,\alpha^k)|t)\kappa(t) \quad > \quad q_j \int_t \alpha_j^k(t)\kappa(dt)$$

for sufficiently large k, which implies $\tilde{z}_j^k D(j,y^k,\alpha^k) z_j^k$, a contradiction. □

Appendix C. Lower Measurability of Extreme Points

This section of the appendix contains a lemma establishing lower measurability of the closure of extreme points of a correspondence. Note that by Lemma 18.3 of AB, this result extends to the correspondence $\mathrm{ext}\varphi(\cdot)$ of extreme points as well (although this correspondence may not have closed values).

Lemma C.1. *Let (S,Σ) denote a measurable space, and assume $\varphi: S \rightrightarrows \mathfrak{R}^d$ is lower measurable with nonempty and compact values. Then the correspondence $s \mapsto \overline{ext}\varphi(s)$ is lower measurable with nonempty and compact values.*

Proof. Nonempty and compact values follow from compactness of $\varphi(s)$. To prove lower measurability, let $\psi(s) = \mathrm{co}\varphi(s)$, and note that these sets possess the same extreme points, i.e., $\mathrm{ext}\psi(s) = \mathrm{ext}\varphi(s)$. Thus, it suffices to show that the correspondence $s \mapsto \overline{\mathrm{ext}}\psi(s)$ is lower measurable. Let $\{x^m\}$ be a countable, dense subset of $\mathfrak{R}^d$, and for each m, define the continuous mapping $d_m: \mathfrak{R}^d \to \mathfrak{R}$

by $d_m(x) = ||x^m - x||$. By a measurable version of the maximum theorem (see Theorem 18.19 of AB), the correspondence $\Phi^m : S \rightrightarrows \mathfrak{R}^d$ defined by

$$\Phi^m(s) \quad = \quad \arg\max\{d_m(x) : x \in \psi(s)\}$$

is lower measurable. By Corollary 7.87 of AB, $\Phi^m(s)$ is contained among the exposed points of $\psi(s)$, and therefore $\Phi^m(s) \subseteq \text{ext}\psi(s)$.[22] By Lemma 18.4 of AB, it follows that the correspondence $\Phi : S \rightrightarrows \mathfrak{R}^d$ defined by $\Phi(s) = \bigcup_{m=1}^{\infty} \Phi^m(s)$ is lower measurable.

Given any $s \in S$, we claim that $\Phi(s)$ is dense among the exposed points of $\psi(s)$. Let y be any exposed point of $\psi(s)$, and let $f : \mathfrak{R}^d \to \mathfrak{R}$ be a linear function such that $\arg\max\{f(x) : x \in \psi(s)\} = \{y\}$, i.e., letting p be the gradient of f normalized so that $||p|| = 1$, we have $p \cdot x < p \cdot y$ for all $x \in \psi(s)$ with $x \neq y$. Consider any $\epsilon > 0$, and define $z_n = y - np$. We will prove that for $n > 0$ large enough, we have $\arg\max\{||z_n - x|| : x \in \psi(s)\} \subseteq B_\epsilon(y)$. If not, then for arbitrarily large n, there exists $v_n \in \psi(s) \setminus B_\epsilon(y)$ such that $||z_n - v_n|| \geq ||z_n - y|| = n$. By compactness of $\psi(s)$, we may assume $v_n \to v \in \psi(s)$. Since y uniquely maximizes f on $\psi(s)$ and $y \neq v$, there exists $a > 0$ such that $p \cdot v + a < p \cdot y$. Setting $w = y - ap$, we have $p \cdot v = p \cdot w$, and in particular, the vectors $v - w$ and $w - z_n$ are orthogonal. It follows that

$$||z_n - v|| \; = \; \sqrt{||z_n - w||^2 + ||v - w||^2} \; = \; \sqrt{(n-a)^2 + ||v - w||^2},$$

which is strictly less than n for n great enough. This implies $v \in B_n(z_n)$ for high enough n. Furthermore, the sequence $\{B_n(z_n)\}$ is increasing in the sense of set inclusion, for given any $x \in B_n(z_n)$, we have $||z_{n+1} - x|| \leq ||z_{n+1} - z_n|| + ||z_n - x|| \leq n + 1$, implying $B_n(z_n) \subseteq B_{n+1}(z_{n+1})$. We conclude that $||z_n - v_n|| < n$ for high enough n, a contradiction. Thus, $\arg\max\{||z_n - x|| : x \in \psi(s)\} \subseteq B_\epsilon(y)$ for some n. Since $\{x^m\}$ is dense in $\mathfrak{R}^d$, we may approximate z_n to an arbitrary degree by elements x^m, and then the theorem of the maximum (Theorem 17.31 of AB) implies that $\Phi^m(s) \subseteq B_\epsilon(y)$ for some m, and therefore $\Phi(s) \cap B_\epsilon(s) \neq \varnothing$. We conclude that $\Phi(s)$ is dense among the exposed points of $\psi(s)$, as claimed.

Finally, Theorem 7.89 of AB implies that the exposed points of $\psi(s)$ are dense among the extreme points of $\psi(s)$, and thus $\overline{\Phi}(s) = \overline{\text{ext}}\psi(s)$ for all $s \in S$. Then lower measurability of $s \mapsto \overline{\text{ext}}\psi(s)$ follows from Lemma 18.3 of AB. $\square$

[22]Given a set $A \subseteq \mathfrak{R}^d$, we say $x \in \mathfrak{R}^d$ is a *strongly exposed point* of A if it is the unique maximizer over A of a linear function.

References

[1] R. Agarwal and D. O'Regan (2002) "Fixed-Point Theory for Weakly Sequentially Upper-Semicontinuous Maps with Applications to Differential Inclusions," *Nonlinear Oscillations,* 5: 277–286.

[2] C. Aliprantis and K. Border (2006) *Infinite Dimensional Analysis: A Hitchhiker's Guide, Third Edition,* New York: Springer.

[3] Z. Artstein (1979) "A Note on Fatou's Lemma in Several Dimensions," *Journal of Mathematical Economics,* 6: 277-282.

[4] Z. Artstein (1989) "Parametrized Integration of Multifunctions with Applications to Control and Optimization," *SIAM Journal of Control and Optimization,* 27: 1369–1380.

[5] J.-P. Aubin and H. Frankowska (1990) *Set-Valued Analysis,* Boston: Birkhauser.

[6] R. Aumann (1965) "Integrals of Set-valued Functions," *Journal of Mathematical Analysis and Applications,* 12: 1–12.

[7] E. Balder (1990) "New Sequential Compactness Results for Spaces of Scalarly Integrable Functions," *Journal of Mathematical Analysis and Applications,* 151: 1-16.

[8] E. Balder (2000) "Incompatibility of the Usual Conditions for Equilibrium Existence in Continuum Economies without Ordered Preferences," *Journal of Economic Theory,* 93: 110-117.

[9] E. Balder (2002) "A Unifying Pair of Cournot-Nash Equilibrium Existence Results," *Journal of Economic Theory,* 102: 437–470.

[10] P. Barelli and J. Duggan (2014) "Purification of Bayes Nash Equilibrium with Correlated Types and Interdependent Payoffs," unpublished, University of Rochester.

[11] A. Caplin and B. Nalebuff (1997) "Competition among Institutions," *Journal of Economic Theory,* 72: 306–342.

[12] G. Carmona and K. Podczeck (2009) "Existence of Pure-Strategy Equilibria in Large Games," *Journal of Economic Theory,* 144: 1300–1319.

[13] G. Carmona and K. Podczeck (2014) "Existence of Nash Equilibrium in Games with a Measure Space of Players and Discontinuous Payoff Functions," *Journal of Economic Theory,* 152: 130–178.

[14] J. Duggan (2011) "General Conditions for the Existence of Maximal Elements via the Uncovered Set," *Journal of Mathematical Economics,* 47: 755–759.

[15] J. Duggan (2012a) "Noisy Stochastic Games," *Econometrica,* 80: 2017–2045.

[16] J. Duggan (2012b) "Measurable Maximum Lemma for Non-ordered Preferences," unpublished, University of Rochester.

[17] J. Duggan and T. Kalandrakis (2012) "Dynamic Legislative Policy Making," *Journal of Economic Theory,* 147: 1653–1688.

[18] D. Epple, R. Filimon, and T. Romer (1984) "Equilibrium among Local Jurisdictions: Toward an Integrated Treatment of Voting and Residential Choice," *Journal of Public Economics,* 24: 281–308.

[19] J. Greenberg and B. Shitovitz (1988) "Consistent Voting Rules for Competitive Local Public Goods Economies," *Journal of Economic Theory,* 46: 223–236.

[20] M. Greinecker and K. Podczeck (2014) "Purification and Roulette Wheels," *Economic Theory*, forthcoming.

[21] W. Hildenbrand (1974) *Core and Equilibria of a Large Economy,* Princeton: Princeton University Press.

[22] A. Jaśkiewicz and A. S. Nowak (2014) "Stationary Markov Perfect Equilibria in Risk-Sensitive Stochastic Overlapping Generations Models," *Journal of Economic Theory,* forthcoming.

[23] M. A. Khan, K. Rath, and Y. Sun (1997) "On the Existence of Pure Strategy Equilibria in Games with a Continuum of Players", *Journal of Economic Theory,* 76: 13–46.

[24] M. A. Khan, K. Rath, Y. Sun, and H. Yu (2013) "Large Games with a Bio-Social Typology," *Journal of Economic Theory*, forthcoming.

[25] H. Konishi (1996) "Voting with Ballots and Feet: Existence of Equilibrium in a Local Public Good Economy," *Journal of Economic Theory,* 68: 480–509.

[26] V. Martins-da-Rocha and M. Topuzu (2008) "Cournot-Nash Equilibria in Continuum Games with non-Ordered Preferences," *Journal of Economic Theory,* 140: 314–327.

[27] A. Mas-Colell (1984) "On a Theorem of Schmeidler," *Journal of Mathematical Economics,* 13: 201–206.

[28] M. Noguchi (2009) "Existence of Nash Equilibria in Large Games," *Journal of Mathematical Economics,* 45: 168–184.

[29] A. Nowak and T. Raghavan (1992) "Existence of Stationary Correlated Equilibria with Symmetric Information for Discounted Stochastic Games," *Mathematics of Operations Research,* 17: 519–526.

[30] K. Podczeck (2009) "On Purification of Measure-valued Maps," *Economic Theory,* 38: 399–418.

[31] M. Rao (2005) *Conditional Measures and Applications, Second Edition* New York: CRC Press.

[32] K. Rath (1992) "A Direct Proof of the Existence of Pure Strategy Equilibria in Games with a Continuum of Players," *Economic Theory,* 2: 427–433.

[33] D. Schmeidler (1973) "Equilibrium Points of Nonatomic Games," *Journal of Statistical Physics,* 7: 295–300.

[34] C. Tiebout (1956) "A Pure Theory of Local Expenditures," *Journal of Political Economy,* 64: 416–424.

[35] F. Westhoff (1977) "Existence of Equilibria in Economies with a Local Public Good," *Journal of Economic Theory,* 14: 84–112.

[36] N. Yannelis (1991) "Integration of Banach-Valued Correspondences," unpublished, University of Illinois.

[37] H. Yu and W. Zhu (2005) "Large Games with Transformed Summary Statistics," *Economic Theory,* 26: 237–241.