Games and Economic Behavior 93 (2015) 42–58

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Markets, correlation, and regret-matching ☆

Sergiu Hart [a,*], Andreu Mas-Colell [b]

[a] Federmann Center for the Study of Rationality, Department of Economics, and Einstein Institute of Mathematics, The Hebrew University of Jerusalem, Feldman Building, Givat Ram, 9190401 Jerusalem, Israel

[b] Department of Economics and Business, Universitat Pompeu Fabra, and Barcelona Graduate School of Economics, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain

ARTICLE INFO

Article history:
Received 17 October 2013
Available online 15 July 2015

JEL classification:
C7
D5

Keywords:
Exchange economies
Walrasian equilibrium
Correlated equilibrium
Sunspot equilibrium
Dynamics
Regret-matching
Socially concave games
Deterministic regret-matching dynamics
Hannan set
Implementation of Walrasian equilibria

ABSTRACT

Inspired by the existing work on correlated equilibria and regret-based dynamics in games, we carry out a first exploration of the links between equilibria and dynamics in (exchange) economies. The leading equilibrium concept is Walrasian equilibrium, and the dynamics (specifically, regret-matching dynamics) apply to trading games that fit the economic structure and whose pure Nash equilibria implement the Walrasian outcomes. Interestingly, in the case of quasilinear utilities (or "transferable utility"), all the concepts essentially coincide, and we get simple deterministic dynamics converging to Walrasian outcomes. Connections to sunspot equilibria are also studied.

1. Introduction

In this paper, inspired by previous work on correlated equilibria and regret-based dynamics in games (see the book of Hart and Mas-Colell, 2013), we carry out a first exploration of the link between the leading equilibrium concept for (exchange) economies, Walrasian equilibrium, and dynamics. We consider no-regret dynamics in trading games that fit the economic structure and have the property that their pure Nash equilibria implement the Walrasian outcomes.

After describing in Section 2 the standard economic model and equilibrium concept, in Section 3 we present the economic game we will focus on. It can be thought of as a very stylized representation of the underlying trading mechanism that generates the economic outcomes. It involves distinguishing one commodity as a means of payment, and one agent as a price controller ("market-maker"). Needless to say, we cannot claim that the game is in any way canonical, but it does

☆ Previous versions: December 2009; September 2013 (Center for Rationality DP-648), September 2014. We thank Yakov Babichenko, Elchanan Ben-Porath, Eddie Dekel, Françoise Forges, Gil Kalai, Yishay Mansour, Yusufcan Masatlioglu, Uri Nadav, and Noam Nisan for useful comments and discussions. We also thank the anonymous referees for their very careful reading of the paper and their insightful suggestions. Research of the first author was partially supported by Advanced Investigator Grant 249159 of the European Research Council (ERC).

* Corresponding author.
E-mail addresses: hart@huji.ac.il (S. Hart), andreu.mas-colell@upf.edu (A. Mas-Colell).
URL: http://www.ma.huji.ac.il/hart (S. Hart).

http://dx.doi.org/10.1016/j.geb.2015.06.009
0899-8256/

capture in a simple way some of the essential features that need to be taken into account: feasibility (i.e., demand equals supply) out of equilibrium, price-taking behavior, and prices controlled by players in the game (i.e., no artificial extraneous players). And we show that, indeed, its pure Nash equilibria correspond to the Walrasian outcomes.

In Section 4 we proceed to analyze the correlated equilibria of the economic game for general economies. Correlated equilibrium (Aumann, 1974) is the generalization of Nash equilibrium that allows for the possibility of noncooperative coordination through (payoff-irrelevant) signaling devices, something that, as the extensive literature on sunspot equilibrium (see the review in Shell, 2007) has persuasively argued, is quite relevant and merits careful consideration in economics. Our results in this section, however, are mostly negative. We show by examples that, beyond straightforward cases, the correlated equilibria of the game, which may be abundant, do not correspond with natural notions of sunspot equilibria of the economy. The difficulty lies in the asymmetric positioning, absent in the specification of the economy, of the distinguished agent (the market-maker) in the game.

The picture changes drastically if utilities are quasilinear (and so utility is transferable across traders) and the distinguished commodity of the game is the economic numeraire. In Section 7 we show that the set of outcomes generated by the correlated equilibria of the game (and even those associated with the more general concept of Hannan equilibria) coincides with the set of Walrasian outcomes, and that the same is true, appropriately specified, for the sunspot outcomes. We proceed then to analyze regret-based dynamics, by defining deterministic (unconditional-)regret-matching ("DURM") strategies, and showing that they guarantee, if played by all players, convergence of outcomes to the Walrasian outcomes.

Sections 5 and 6 develop game-theoretic tools that yield, as an application, the results of Section 7. Section 5 deals with the class of "socially concave games," introduced by Even-Dar et al. (2009), where correlated equilibria (and Hannan equilibria) yield on average the same outcomes as pure Nash equilibria. Section 6 develops deterministic regret-matching dynamics for games with convex action spaces (as is the case for the economic game), and studies their limit "no-regret" behavior.

2. The economic model

There is a finite set $N = \{1, 2, \ldots, n\}$ of *agents* (or *traders*) and a finite set $M = \{1, 2, \ldots, m\}$ of *commodities*. The consumption space of each agent $i \in N$ is[1] $\mathbb{R}^M$, over which he has a utility function $u^i : \mathbb{R}^M \to \mathbb{R}$. We assume the following for every agent $i \in N$:

- The utility function u^i is concave (and thus continuous).
- The marginal rates of substitution are uniformly bounded: there is a finite $K \geq 1$ such that at every $x \in \mathbb{R}^M$, if $0 \neq p \in \mathbb{R}^M$ are supporting prices at x (i.e., $u^i(x') \geq u^i(x)$ implies $p \cdot x' \geq p \cdot x$), then[2] $p \gg 0$ and $1/K \leq p_\ell / p_{\ell'} \leq K$ (and thus u^i is strictly monotonic).

Without loss of generality we may well take the initial endowment of each $i \in N$ to be $0 \in \mathbb{R}^M$, and so a consumption vector $x^i \in \mathbb{R}^M$ is in fact a net trade vector. Let $\mathcal{E}$ denote the economy.

A *Walrasian equilibrium* (or *competitive equilibrium*) of $\mathcal{E}$ consists of a price vector $p \in \mathbb{R}^M_{++}$ and an allocation $(x^i)_{i \in N}$ with $x^i \in \mathbb{R}^M$ and $\sum_{i \in N} x^i = 0$ such that x^i is a demand of agent i at p, for every $i \in N$ (i.e., $p \cdot x^i = 0$, and $p \cdot x' \leq 0$ implies $u^i(x') \leq u^i(x^i)$). It will be convenient to normalize the price vectors so that $p_m = 1$ (where m is the last coordinate of M), and thus $\tilde{p} \in P := \{p \in \mathbb{R}^M_{++} : p_m = 1 \text{ and } 1/K \leq p_\ell / p_{\ell'} \leq K \text{ for all } \ell, \ell' \in M\}$.

Throughout this paper we will refer to $(x^1, \ldots, x^n, p) \in \mathbb{R}^M \times \ldots \times \mathbb{R}^M \times P$ with $\sum_{i \in N} x^i = 0$, i.e., an allocation of the goods together with a price vector, as an ***outcome***. Let WEO denote the set of Walrasian Equilibrium Outcomes.

3. A game implementing the Walrasian equilibria

We provide a simple and natural game whose pure Nash equilibria correspond to the Walrasian equilibria of the market. Here, "corresponds" means that the outcomes $(x^1, \ldots, x^n, p) \in \mathbb{R}^M \times \ldots \times \mathbb{R}^M \times P$ are the same. We impose two restrictions on ourselves. First, there should be no additional artificial agents ("referees" or "designers"); only the given n economic agents play. And second, the game should be "playable," in the sense that the outcome should always be feasible, out of equilibrium as well as at equilibrium.

The game $\mathcal{G}$ generated from the given economy $\mathcal{E}$ is constructed as follows. We single out one agent, say agent n—call him the "*market-maker*"—and one commodity, say commodity m—call it the "*numeraire*." An action of agent $i \neq n$ is a consumption vector of the nonnumeraire goods[3] $y^i \in \mathbb{R}^{M-1}$, and an action of agent $i = n$ is a price vector for the nonnumeraire goods $q \in Q$, where Q is the projection of P to its first $M - 1$ coordinates (i.e., $Q := \{q \in \mathbb{R}^{M-1}_{++} : 1/K \leq q_\ell \leq K \text{ and } 1/K \leq q_\ell / q_{\ell'} \leq K \text{ for all } \ell, \ell' \in M \backslash \{m\}\}$). An n-tuple of actions $a := (y^1, \ldots, y^{n-1}, q)$ generates the outcome $\theta := (x^1, \ldots, x^n; p)$ given by

[1] This assumption, by abstracting away from issues such as nonnegativity constraints on consumption, allows for simple and general trading mechanisms. It may be viewed as replacing stringent constraints with large but finite penalties, and it leads to a more transparent analysis. See however Sections 3.1 and 3.2, where we show ways to deal with such constraints.

[2] $p \gg 0$ means $p_\ell > 0$ for all ℓ, i.e., $p \in \mathbb{R}^M_{++}$.

[3] We write $M - 1$ rather than the correct $M \backslash \{m\}$.

$$y^n := -\sum_{i \neq n} y^i,$$
$$x^i := (y^i, -q \cdot y^i) \in \mathbb{R}^M \text{ for all } i = 1, 2, \ldots, n, \tag{1}$$
$$p := (q, 1) \in P.$$

Thus $\sum_{i \in N} x^i = 0$ and $p \cdot x^i = 0$ for all $i \in N$, and so the allocation of goods outcome is always feasible (this is the "playability" condition), and every agent's consumption lies on his budget line for the price vector p. Finally, the payoff functions are defined as[4]

$$g^i(y^1, \ldots, y^{n-1}, q) := u^i(x^i) \quad \text{for } i \neq n, \text{ and}$$
$$g^n(y^1, \ldots, y^{n-1}, q) := u^n(x^n) - \sup\{u^n(z) : z \in \mathbb{R}^M \text{ and } p \cdot z = 0\}.$$

The interpretation is as follows. Each agent $i \neq n$ chooses his consumption of the nonnumeraire goods, and the market-maker n chooses the prices. The market-maker's consumption of the regular goods is determined by the other agents: he absorbs all the excess demand or supply.[5] Every agent "pays" for his consumption with the numeraire good, according to the prices determined by the market-maker. The objective of each agent $i \neq n$ is to maximize his utility of his consumption of all goods (including the numeraire); and that of the market-maker n is to minimize the difference between the maximal utility he could have obtained at the prices that he chose and the utility of his actual consumption. In other words, he wants to minimize his "dissatisfaction" over his consumption not being optimal at the prices he chose.

Remarks. *(a)* Game-theoretic discipline imposes, on the one hand, that the allocations resulting from arbitrary play should be feasible, and, on the other, that prices should be determined strategically within the game. These two requirements do in fact reinforce each other. Because the consumptions of all the different agents cannot be determined independently, we may free one agent from the task. This makes him available for the determination of prices.

(b) In their classical existence proof, Arrow and Debreu (1954) postulated a virtual "referee" in charge of prices. This is not unlike what we do. The main difference is that we do not use an extra agent: our market-maker, who determines the prices, is one of the agents in the economy. This accounts for the difference between the objective of the Arrow–Debreu referee (maximize the value of the excess demand) and that of our market-maker (given by g^n).

(c) We have adopted an approach that treats the agents asymmetrically: the entire weight of adjusting prices is borne by one of the agents. In this sense our approach is in the spirit of the price-adjustment processes led by a referee. But this is not the only possible way. Games that are symmetric among the traders and implement Walrasian outcomes have been constructed, starting with Hurwicz (1979) and Schmeidler (1980).[6] It would certainly be of interest to analyze the issues we study here (and in particular the dynamics of Section 6) for these trading games.

The equilibrium concept that we consider for the game $\mathcal{G}$ is that of a *pure Nash equilibrium*. Let PNEO stand for the set of Pure Nash Equilibrium Outcomes of $\mathcal{G}$, i.e., the set of outcomes generated (see (1)) by pure Nash equilibria of the game $\mathcal{G}$.

Theorem 1. *The set of pure Nash equilibrium outcomes of the game $\mathcal{G}$ coincides with the set of Walrasian equilibrium outcomes of the economy $\mathcal{E}$:* PNEO = WEO.

This readily follows from the lemma below that deals with the special agent n. Note first that $g^n(y^1, \ldots, y^{n-1}, q) = u^n(y^n, -q \cdot y^n) - \sup_{w \in \mathbb{R}^{L-1}} u^n(w, -q \cdot w) \leq 0$; we will show that, for every y^n, the market-maker n can always choose q in Q so that his payoff is maximal, i.e., 0, and so his consumption $x^n = (y^n, -q \cdot y^n)$ is his demand at the price vector $p = (q, 1)$.

Lemma 2. *For every $y \in \mathbb{R}^{M-1}$ there exists $q \in Q$ such that*

$$u^n(y, -q \cdot y) = \sup_{w \in \mathbb{R}^{M-1}} u^n(w, -q \cdot w).$$

Proof. For each real α, put $z(\alpha) := (y, \alpha)$, and let $P(\alpha) \subset P$ be the set of supporting prices for agent n at the point $z(\alpha)$; i.e., $p \in P(\alpha)$ if and only if $p \cdot x \geq p \cdot z(\alpha)$ for any x in the convex set $\{x \in \mathbb{R}^M : u^n(x) \geq u^n(z(\alpha))\}$. Clearly $P(\alpha)$ is a nonempty convex compact set; moreover, the correspondence $P(\cdot)$ is upper-hemicontinuous.[7] Put $F(\alpha) := \{p \cdot z(\alpha) : p \in P(\alpha)\} \subset \mathbb{R}$;

[4] It may well happen that the payoff of agent n is $-\infty$. To avoid this, one may restrict the supremum to vectors z that lie in some bounded convex set containing x^n in its interior (see the Proof of Lemma 2 below). Also, see the Remark following the Proof of Theorem 1 for other possible payoff functions g^n.

[5] Our market-maker is quite similar to Nasdaq market-makers.

[6] See also Groves and Ledyard (1987), Hurwicz et al. (1995), and Postlewaite and Wettstein (1983).

[7] This is standard. Take $\alpha_m \to \alpha$ and $p_m \to p$ with $p_m \in P(\alpha_m)$. Then $u^n(x) \geq u^n(z(\alpha))$ implies $u^n(x + (\varepsilon, \ldots, \varepsilon)) > u^n(z(\alpha_m))$ for $\varepsilon > 0$ and m large enough, and hence $p_m \cdot (x + (\varepsilon, \ldots, \varepsilon)) \geq p_m \cdot z(\alpha_m)$, which in the limit yields $p \cdot x \geq p \cdot z(\alpha)$.

then F is also a convex-valued upper-hemicontinuous correspondence. Since $p \in P(\alpha)$ implies $1/K \leq p_\ell \leq K$, it follows that all elements of $F(\alpha)$ are negative for small enough α, and positive for large enough α; therefore there is α^* such that $0 \in F(\alpha^*)$, which means that there is $p^* = (q^*, 1) \in P$ such that $0 = p^* \cdot z(\alpha^*) = q^* \cdot y + \alpha^*$, or $\alpha^* = -q^* \cdot y$. Since p^* is a supporting price at $z(\alpha^*) = (y, -q^* \cdot y)$, the result follows. Indeed, if there is w such that $u^n(w, -q^* \cdot w) > u^n(y, -q^* \cdot y)$, then $u^n(w, -q^* \cdot w - \varepsilon) > u^n(y, -q^* \cdot y)$ for small enough $\varepsilon > 0$, and then $p^* \cdot (w, -q^* \cdot w - \varepsilon) = -\varepsilon < 0 = p^* \cdot z(\alpha^*)$, a contradiction to $p^* \in P(\alpha^*)$. $\square$

Proof of Theorem 1. Let $(\tilde{y}^1, \ldots, \tilde{y}^{n-1}, \tilde{q})$ be a pure Nash equilibrium, with corresponding outcome $(\tilde{x}^1, \ldots, \tilde{x}^n, \tilde{p})$. For each agent $i \neq n$, the bundle $\tilde{x}^i$ satisfies $\tilde{p} \cdot \tilde{x}^i = 0$ (by definition of $\tilde{x}^i$) and maximizes g^i over y^i; hence it maximizes u^i over x^i in i's budget set for the prices $\tilde{p}$. For agent n, the same conclusion is reached using Lemma 2. Altogether, the outcome is indeed a Walrasian equilibrium outcome.

Conversely, given a Walrasian equilibrium outcome $(\tilde{x}^1, \ldots, \tilde{x}^n, \tilde{p})$, the corresponding actions $(\tilde{y}^1, \ldots, \tilde{y}^{n-1}, \tilde{q})$ clearly yield a pure Nash equilibrium of the game (for $i = n$, use Lemma 2). $\square$

Remark. One could replace g^n with any other payoff function that satisfies Lemma 2, i.e., the payoff is 0 whenever x^n is a demand of n at p, and it is < 0 otherwise (for instance, take the minimal expenditure needed for agent n to reach utility level $u^n(x^n)$ when the price vector is p).

3.1. Compact action spaces

While the action space Q of agent n is compact, those of the other $n-1$ agents are not. It makes sense, and it will be important for the dynamic analysis of later sections, to ask if the result of Theorem 1 also holds when the action sets of these agents are restricted to a suitable compact box, say $B := [-b, b]^{M-1}$. Call this game $\mathcal{G}_b$. Without making further assumptions on the economy, the answer is negative. As an example, take $n = 2$ and $m = 2$, and $u^1(x) = 2x_1 + x_2$, and $u^2(x) = x_1 + x_2$. It is immediate to see that this economy has no Walrasian equilibrium, whereas every $\mathcal{G}_b$ possesses a pure Nash equilibrium (the payoff functions are strategically equivalent[8] to concave functions and the action sets are convex compact sets). This suggests that it is necessary to resort to some assumption of the type that guarantees the existence of Walrasian equilibria.

One simple way to do so is as follows (cf.[9] O. Hart, 1974). For each agent i, a vector $z \in \mathbb{R}^M$ is an (improving) *recession direction* if $u^i(\lambda z) \geq u^i(0)$ for every[10] $\lambda \geq 0$; let $C^i \subset \mathbb{R}^M$ be the closed cone of all recession directions of agent i. Recall that a closed cone C is *pointed* if $C \backslash \{0\}$ is *strictly* included in a half-space; equivalently,[11] $c_1, \ldots, c_J \in C$ satisfy $\sum_{j=1}^J c_j = 0$ if and only if $c_1 = \ldots = c_J = 0$. The assumption is:

(C) There exists a *pointed cone* C containing the recession directions of all agents, i.e., $C^i \subset C$ for all $i \in N$.

Roughly speaking, (C) says that there is some degree of commonality among the agents in the directions of improvement.

An *individually rational allocation* is $(x^1, \ldots, x^n) \in \mathbb{R}^M \times \ldots \times \mathbb{R}^M$ with $\sum_{i=1}^n x^i = 0$ such that $u^i(x^i) \geq u^i(0)$ for all agents $i \in N$. From (C) we readily have:

Lemma 3. *Let $\mathcal{E}$ satisfy (C). Then there exists $b_0 > 0$ such that any individually rational allocation $(x^1, \ldots, x^n)$ satisfies $||x^i|| < b_0$ for all $i \in N$.*

Proof. By contradiction: Let $(x_s^1, \ldots, x_s^n)_{s=1,2,\ldots}$ be a sequence of individually rational allocations with $\mu_s := \sum_{j=1}^n ||x_s^j|| \to_s \infty$. Take an appropriate subsequence so that $(1/\mu_s)x_s^i$ converges for each i, say, to z^i. Then $u^i(\lambda z^i) \geq u^i(0)$ for (a sequence of) arbitrarily large λ, which implies that $z^i \in C^i \subset C$. But $\sum_i z^i = 0$ (since $\sum_i x_s^i = 0$) and not all z^i are 0 (since $||\sum_i z^i|| = 1$), contradicting the pointedness of C. $\square$

Corollary 4. *Let $\mathcal{E}$ satisfy (C). If $(x^1, \ldots, x^n)$ is a Walrasian allocation then $||x^i|| < b_0$ for all i.*

Proof. A Walrasian allocation is individually rational.[12] $\square$

---

[8] Cf. Section 5.

[9] In the context of financial economies, where unbounded consumption sets arise naturally.

[10] These are the recession directions *at* 0; since u^i is a concave function, they are the same at every $x \in \mathbb{R}^M$; see Rockafellar (1966, Theorem 2A; and 1970, Theorem 8.7).

[11] Or, there is a vector d such that $d \cdot c > 0$ for every $c \in C \backslash \{0\}$.

[12] Condition (C) implies existence of Walrasian equilibria. Indeed, restrict the preferences of every agent to a translated orthant that contains the vector $(-b_0, \ldots, -b_0)$, where b_0 is as in Lemma 3, and apply standard proofs (e.g., Debreu, 1959).

Corollary 5. *Let $\mathcal{E}$ satisfy (C). If $(x^1, \ldots, x^n, p)$ is a pure Nash equilibrium outcome of the game $\mathcal{G}$ then $||x^i|| < b_0$ for all i. Moreover, $\mathcal{G}_b$ and $\mathcal{G}$ have the same pure Nash equilibria, for every $b \geq b_0$.*

Proof. The first statement follows from Theorem 1 and Corollary 4.

Next, any pure Nash equilibrium of $\mathcal{G}$ (which, as we have just shown, is feasible in $\mathcal{G}_b$) is clearly also a pure Nash equilibrium of $\mathcal{G}_b$. Conversely, if $(\tilde{y}^1, \ldots, \tilde{y}^{n-1}, \tilde{q})$ is a pure Nash equilibrium of $\mathcal{G}_b$ with outcome $(x^1, \ldots, x^n, p)$, then $u^i(x^i) \geq u^i(0)$ for all $i \leq n-1$ (since choosing $y^i = 0$, which yields $x^i = 0$, is always possible), and also $u^n(x^n) \geq u^n(0)$ (by Lemma 2, x^n is a demand of n at the price vector p). Thus $(x^1, \ldots, x^n)$ is an individually rational allocation, and so, in particular, $||\tilde{y}^i|| \leq ||x^i|| < b_0$ for all $i \leq n-1$ (by Lemma 3). Thus $\tilde{y}^i$ is an interior maximizer of a concave function (specifically, the function $u^i(y^i, -\tilde{q} \cdot y^i)$ of y^i), and so it is in fact a global maximizer; therefore $(\tilde{y}^1, \ldots, \tilde{y}^{n-1}, \tilde{q})$ is a pure Nash equilibrium of $\mathcal{G}$. $\square$

3.2. Nonnegative consumptions

Up to now we have assumed that the net trade x^i of each agent i is unbounded, and therefore so is his consumption. In standard economic models it is however usual to have the consumption bounded from below, specifically, to be nonnegative. We will show here a way to handle such models.

Let $\mathbb{R}^M_+$ be the consumption set (and so $u^i : \mathbb{R}^M_+ \to \mathbb{R}$) and let $e^i \in \mathbb{R}^M$ be the initial endowment of agent[13] i, for $i = 1, \ldots, n$. We make the following commonly used assumptions for every agent $i \in N$.[14]

- Each good $\ell = 1, \ldots, m$ is "indispensable": $u^i(z) > 0$ for all $z \in \mathbb{R}^M_{++}$, and $u^i(z) = 0$ for all $z \in \mathrm{bd}\,\mathbb{R}^M_+$ (i.e., to get a positive utility one needs all goods).[15]
- Each good ℓ is "useful": u^i is strictly increasing in all coordinates on $\mathbb{R}^M_{++}$.
- The initial endowments are strictly positive: $e^i \in \mathbb{R}^M_{++}$ for all $i = 1, \ldots, n$.

Consider the set of individually rational allocations $(z^1, \ldots, z^n)$, i.e., $\sum_i z^i = \sum_i e^i$ and $u^i(z^i) \geq u^i(e^i)$ for each i. The assumptions above imply that each z^i belongs to a compact subset B^i of $\mathbb{R}^M_{++}$ (because the set of individually rational allocations is closed, $0 \leq z^i \leq \sum_i e^i$, and moreover $u^i(z^i) \geq u^i(e^i) > 0$ implies $z^i \in \mathbb{R}^M_{++}$), and so any function $\tilde{u}^i$ that coincides with u^i on B^i yields the same Walrasian equilibrium outcomes (which are individually rational). We may therefore take the functions $\tilde{u}^i$ for all i to be defined on the whole space $\mathbb{R}^M$, and moreover to satisfy the assumptions of Section 2 (specifically, marginal rates of substitution uniformly bounded by some K).[16] Furthermore, there is no need to appeal to the cone condition (C) of Section 3.1, since $b_0 = ||\sum_i e^i||$ serves as a bound on the individually rational, hence equilibrium, outcomes.[17] Finally, we translate the origin for each agent i from e^i to 0.

One way to view the change from u^i to $\tilde{u}^i$ is that crossing the boundary of $\mathbb{R}^M_+$ incurs an infinite penalty in u^i (i.e., it is not allowed; equivalently, $u^i(z) = -\infty$ for $z \notin \mathbb{R}^M_+$), and a finite penalty in $\tilde{u}^i$; but this penalty is large enough so that the individually rational outcomes are not affected—and so our results hold.

Remark. Concerning playability, we do not know of a way to always ensure $z^i \geq 0$ out of equilibrium (this is a known issue; cf. Hurwicz et al., 1995). Our setup allows this constraint to be violated, albeit at a steep cost.

4. Correlated equilibria and sunspot equilibria

In the previous section we showed that for general economies the pure Nash equilibria of the associated game implement the Walrasian equilibria of the economy. This leads one to ask which economic outcomes would be implemented by the generalization of pure Nash equilibria that allows the possibility of random, payoff-irrelevant signals, namely, *correlated equilibria* (Aumann, 1974). Obviously, there may be other outcomes beside Walrasian equilibria. Since payoff-irrelevant signals are involved, the natural candidates are some versions of *sunspot equilibria* (Cass and Shell, 1991; Shell, 2007). These are explored in this section.[18]

---

[13] In this section it is more convenient to deal with consumption vectors z^i, rather than with net trades $x^i = z^i - e^i$ as in the previous sections. The translation between the two setups is immediate.

[14] For instance, Cobb–Douglas utilities satisfy these assumptions.

[15] $\mathrm{bd}\,\mathbb{R}^M_+$ denotes the boundary of $\mathbb{R}^M_+$ (i.e., all $z \in \mathbb{R}^M_+$ with at least one coordinate equal to 0). The specific utility value of zero that separates $\mathbb{R}^M_{++}$ from $\mathrm{bd}\,\mathbb{R}^M_+$ is arbitrary, as adding a constant to the utility does not matter.

[16] Define $\tilde{u}^i$ to be the minimum of (the compact set of) the supporting hyperplanes to the graph of u^i restricted to B^i. To make the slopes of $\tilde{u}^i$ on the boundary of $\mathbb{R}^M_+$ steeper, one may include also hyperplanes such as $h(z) = c_\ell z_\ell$ with large $c_\ell > 0$ (so that h lies above u^i on B^i).

[17] The recession cone of the original utilities u^i is $\mathbb{R}^M_+$, and so for appropriate $\tilde{u}^i$ it is close to $\mathbb{R}^M_+$—and condition (C) is satisfied.

[18] While the Walrasian equilibria and the pure Nash equilibria are *ordinal* concepts—depending only on the preferences and not on the specific utility functions—the mixed Nash equilibria, sunspot equilibria, and correlated equilibria are *cardinal* concepts which depend on the specific utility functions. Since

For the definition of sunspot equilibria that we find natural in our context, we offer some illustrative examples and conclude that there is no simple correspondence between sunspot and correlated equilibrium outcomes for general economies. However, for the special class of economies with quasilinear utilities (or "transferable utility") we will see, in Section 7, that the situation is quite different.

4.1. Correlated equilibria

We start by recalling the concept of correlated equilibrium (Aumann, 1974). Let $G = (N, (A^i)_{i \in N}, (g^i)_{i \in N})$ be a game in strategic form, where N is the set of players, A^i is the set of actions of player i, and $g^i : A \to \mathbb{R}$ is the payoff function of player i, where $A := \Pi_{i \in N} A^i$ is the set of action combinations of all players ("action profiles"). The setup consists of:

- The probability space $(\Omega, \mathcal{F}, P)$ of the states of the world.[19]
- For each player $i \in N$, the signal[20] $\mathbf{s}^i(\omega) \in S^i$ that i gets in state ω (which is player i's information on the state of the world; S^i is an arbitrary set).
- For each player $i \in N$, the action $\mathbf{a}^i(\omega) \in A^i$ that i chooses in state ω, depending on his signal[21] $\mathbf{s}^i(\omega)$.
- All this constitutes a *correlated equilibrium* if each player $i \in N$ maximizes his expected payoff given his information, i.e.,

$$\mathbf{a}^i(\omega) \in \arg\max_{b^i \in A^i} \mathbb{E}\left[g^i(b^i, \mathbf{a}^{-i}) \mid \mathbf{s}^i(\omega)\right] \tag{2}$$

  for all[22] $\omega \in \Omega$.

Equivalently, consider the distribution $\mu \in \Delta(A)$ of the random variable $\mathbf{a} = (\mathbf{a}^i)_{i \in N}$ (with values in A). Then μ is the distribution of a correlated equilibrium if and only if

$$a^i \in \arg\max_{b^i \in A^i} \mathbb{E}_{a^{-i} \sim \mu^{-i}(\cdot|a^i)}\left[g^i(b^i, a^{-i})\right] = \arg\max_{b^i \in A^i} \int_{A^{-i}} g^i(b^i, a^{-i})\, \mathrm{d}\mu(a^{-i} \mid a^i)$$

for every player $i \in N$ and every action $a^i \in A^i$, where $A^{-i} := \Pi_{j \neq i} A^j$ and $\mu^{-i}(\cdot|a^i) \in \Delta(A^{-i})$ is the conditional distribution of μ given a^i.

4.2. Sunspot equilibria

Two general ideas underlie the concept of sunspot equilibrium. First, while the fundamentals of the economy are fixed (and known to all agents), the prices may be random. Second, unlike Walrasian equilibria, these prices are not necessarily fully known to all agents, who may possess only partial information on the prices.

Formally:

- $(\Omega, \mathcal{F}, P)$ is the probability space of the states of the world.[23] In each state of the world $\omega \in \Omega$ there is a *price* vector $\mathbf{p}(\omega) \in P$; let $\mathbf{q}(\omega) \in \mathbb{R}^{M-1}$ denote the price vector for the $m-1$ non-numeraire goods, so that $\mathbf{p}(\omega) = (\mathbf{q}(\omega), 1)$.
- Each agent $i \in N$ gets a "signal" $\mathbf{s}^i(\omega)$ on the state of the world (which is agent i's information on the prices).
- Each agent $i \in N$ chooses his consumption $y^i \in \mathbb{R}^{M-1}$ of the non-numeraire goods. His consumption of the numeraire will be determined by the (possibly unknown) prices: he will have to pay $\mathbf{q}(\omega) \cdot y^i$. Thus $\mathbf{x}^i(\omega) = (y^i, -\mathbf{q}(\omega) \cdot y^i) \in \mathbb{R}^M$ will be his consumption vector in state ω. The *demand* of i in state ω, which we denote by $\mathbf{y}^i(\omega)$, is determined—as a function of his information[24] $\mathbf{s}^i(\omega)$—so as to maximize his expected utility given his information, i.e.,

$$\mathbf{y}^i(\omega) \in \arg\max_{y^i \in \mathbb{R}^{M-1}} \mathbb{E}\left[u^i(y^i, -\mathbf{q}(\omega) \cdot y^i) \mid \mathbf{s}^i(\omega)\right] \tag{3}$$

  for all $\omega \in \Omega$.

---

these latter concepts use randomizations, the agents' utility functions are now viewed as von Neumann–Morgenstern utilities that express risk preferences (and so a randomization is evaluated by its expected utility).

[19] A common prior is thus assumed.

[20] Functions of the state of the world ω, i.e., random variables on Ω, appear in bold type.

[21] Without loss of generality one may assume that (i) $\mathbf{a}^i$ is determined by $\mathbf{s}^i$ (formally: $\mathbf{a}^i$ is measurable with respect to $\mathbf{s}^i$)—otherwise, one may refine the signal $\mathbf{s}^i$—and that (ii) $\mathbf{a}^i = \mathbf{s}^i$—otherwise, one may merge signals s^i after which the decision a^i is the same. This is the standard way of presenting a correlated equilibrium.

[22] All probabilistic statements, including those involving conditioning, should be understood to hold almost everywhere (a.e.).

[23] Again, a common prior is assumed.

[24] Thus $\mathbf{y}^i$ is without loss of generality measurable with respect to $\mathbf{s}^i$ (refine $\mathbf{s}^i$ otherwise).

- All this constitutes a *sunspot equilibrium* if the market clears, i.e., total demand equals total supply, in all states of the world:

$$\sum_{i \in N} \mathbf{x}^i(\omega) = 0$$

for all $\omega \in \Omega$ (this is equivalent to $\sum_{i \in N} \mathbf{y}^i(\omega) = 0$).

Thus, in every state of the world total demand equals total supply—i.e., $\sum_{i \in N} \mathbf{x}^i = 0$—and every agent's consumption lies on his budget line—i.e., $\mathbf{p} \cdot \mathbf{x}^i = 0$ for all i.

A Walrasian equilibrium obtains when there is no uncertainty, i.e., when there is a unique state of the world ω (which is therefore commonly known, and so the prices are also commonly known). When all agents have full information (i.e., $\mathbf{s}^i(\omega) = \omega$ for all ω and all i), the price vector $\mathbf{p}(\omega)$ is commonly known in each state of the world ω, and so we get a Walrasian equilibrium in each state ω. This sunspot equilibrium therefore yields a convex combination of Walrasian equilibria.

As we will see below, in general there exist other sunspot equilibria besides the (convex combinations of) Walrasian equilibria; we refer to them as "nontrivial" sunspot equilibria.

Remark. The notion of sunspot equilibrium we use is entirely in the spirit of the very extensive literature on the topic (Cass and Shell, 1991; Shell, 2007). In contrast to the latter, which typically considers multi-period economies, we stick to a static framework, and adapt the definitions accordingly. Issues of correlation have also been discussed in this literature (see, for example, Maskin and Tirole, 1987; Forges and Peck, 1995, and Polemarchakis and Ray, 2006).

4.3. Comparing correlated equilibria and sunspot equilibria

Comparing the sunspot equilibrium condition (3) with the correlated equilibrium condition (2) for the associated game, we immediately see that, for every agent $i \neq n$ (i.e., all agents except the market-maker), the two conditions coincide: both amount to choosing $y^i \in \mathbb{R}^{M-1}$ (i.e., the consumption of the non-numeraire goods) so as to maximize the expected utility of i given the corresponding conditional distribution of prices. However, the two conditions differ importantly for the market-maker n: while in a sunspot equilibrium he behaves exactly like the other agents—choosing his non-numeraire consumption y^n optimally—in a correlated equilibrium his consumption of the non-numeraire goods is determined by the other agents. Instead, he chooses the prices, so as to maximize a suitably specified payoff.

The reason for this difference is that the game must be playable out of equilibrium (a concern that sunspot equilibria do not have), which means that the outcome must be feasible for any combination of actions of all agents. This is resolved, in our approach, by having the consumption of one of the agents (the market-maker) be determined by the consumptions of the other agents.

The examples below will clarify the contrasts between the two concepts.

4.4. Examples

All the examples here have $n = 2$ agents and $m = 2$ goods; in the associated game $\mathcal{G}$, agent 2 serves as "market-maker" and good 2 serves as "numeraire."

Example 6. *A mixed Nash equilibrium that yields a nontrivial*[25] *sunspot equilibrium.*

We seek a Nash equilibrium of the associated game $\mathcal{G}$ where agent 1 plays a pure action, say y', and agent 2, the market-maker, plays a mixed action, say q' and q'', with equal probabilities of 1/2 each.

The allocation of the non-numeraire good (i.e., good 1) is thus $y^1 = y'$ for agent 1 and $y^2 = -y'$ for agent 2. The Nash equilibrium conditions are[26]:

$$\begin{aligned} y' &\in \arg\max_y \frac{1}{2} g^1(y, q') + \frac{1}{2} g^1(y, q''), \\ q' &\in \arg\max_{q \geq 0} g^2(y', q), \\ q'' &\in \arg\max_{q \geq 0} g^2(y', q). \end{aligned}$$

Equivalently (for the last two conditions use Lemma 2):

[25] I.e., a sunspot equilibrium that is not a convex combination of Walrasian equilibria.

[26] In all the examples here the maximizer is in fact unique, and so $y' = \arg\max_y$, etc.

$$y' \in \arg\max_y \frac{1}{2}u^1(y, -q'y) + \frac{1}{2}u^1(y, -q''y),$$

$$-y' \in \arg\max_y u^2(y, -q'y), \tag{4}$$

$$-y' \in \arg\max_y u^2(y, -q''y). \tag{5}$$

These are precisely the conditions for a sunspot equilibrium where agent 1 does not know the price (which has equal probability of being $q' = 1$ or $q'' = 2$), whereas agent 2 knows the price (either q' or q'').

In addition, we want neither $(y', -y', q')$ nor $(y', -y', q'')$ to be a pure Nash equilibrium, so that neither outcome will be Walrasian; thus

$$y' \notin \arg\max_y u^1(y, -q'y),$$

$$y' \notin \arg\max_y u^1(y, -q''y).$$

It is easy to find such an example. Take, for instance, $y' = 10$, $q' = 1$, $q'' = 2$, and standard utility functions such as

$$u^1(x_1, x_2) = 5\log(x_1) + 4\log(x_2 + 30),$$

$$u^2(x_1, x_2) = 2\log(x_1 + 30) + \log x_2,$$

in the regions, say, $(5, -25) \le x \le (15, -5)$ for u^1, and $(-15, 5) \le x \le (-5, 25)$ for u^2; then extend these utility functions to $\mathbb{R}^2$ so that they are concave and satisfy the condition of uniformly bounded marginal rates of substitution. The resulting consumption vectors, which are $x^1 = (10, -10)$ and $(10, -20)$ for agent 1 and $x^2 = (-10, 10)$ and $(-10, 20)$ for agent 2, belong to the above regions. □

Remark. This example illustrates a general result: a Nash equilibrium where the market-maker plays a strictly mixed action is always a sunspot equilibrium, which is moreover nontrivial if the Nash equilibrium is not a mixture of pure Nash equilibria.

Example 7. *A sunspot equilibrium that cannot be obtained from a correlated equilibrium.*

The setup is as in the previous example, but now we require that $y^2 = -y'$ satisfy neither (4) nor (5), but instead satisfy

$$-y' \in \arg\max_y \frac{1}{2}u^2(y, -q'y) + \frac{1}{2}u^2(y, -q''y).$$

This yields a sunspot equilibrium where neither agent knows the price q, but only that it equals q' or q'' with equal probabilities of 1/2 each: indeed, y^i maximizes $(1/2)u^i(y, -q'y) + (1/2)u^i(y, -q''y)$, for both agents $i = 1, 2$. However, if this outcome were obtained in a correlated equilibrium of the associated game, then the market-maker $i = 2$ would know q (as it is his action), but $y^2 = -y'$ (which is constant and thus known) does *not* maximize either $u^2(y^2, -q'y^2)$ or $u^2(y^2, -q''y^2)$. The same holds for agent $i = 1$, and so neither agent can serve as market-maker.

Specifically, we again use $y' = 10$, $q' = 1$, and $q'' = 2$, and take, for instance,

$$u^1(x_1, x_2) = 5\log(x_1) + 4\log(x_2 + 30),$$

$$u^2(x_1, x_2) = 7\log(x_1 + 30) + 6\log(x_2 + 10),$$

in the regions $(5, -25) \le x \le (15, -5)$ for u^1, and $(-15, 5) \le x \le (-5, 25)$ for u^2, appropriately extended to $\mathbb{R}^2$. □

Example 8. *A correlated equilibrium that does not yield a sunspot equilibrium.*

Here we let y take two values y' and y'', and q take two values q' and q''. We construct a correlated equilibrium that puts equal probability of 1/3 on each one of the three points (y', q'), (y', q''), and (y'', q'). The correlated equilibrium conditions are

$$y' \in \arg\max_y \left(\frac{1}{2}g^1(y, q') + \frac{1}{2}g^1(y, q'')\right),$$

$$y'' \in \arg\max_y g^1(y, q'),$$

$$q' \in \arg\max_{q \ge 0} \left(\frac{1}{2}g^2(y', q) + \frac{1}{2}g^2(y'', q)\right),$$

$$q'' \in \arg\max_{q \ge 0} g^2(y', q),$$

where $g^1(y,q) = u^1(y, -qy)$ and $g^2(y,q) = u^2(-y, qy) - \sup_w u^2(-w, qw)$ (recall Lemma 2). Moreover, in order not to get a sunspot equilibrium, we require that agent 2 (the market-maker) not get his demand, i.e.,

$$-y' \notin \arg\max_y \left(\frac{1}{2}u^2(y, -q'y) + \frac{1}{2}u^2(y, -q''y)\right), \tag{6}$$

$$-y'' \notin \arg\max_y u^2(y, -q'y). \tag{7}$$

Indeed, the signal of agent 2 in a sunspot equilibrium would have to distinguish between the two states (y', q'), (y', q'') (where his demand is $-y'$) and the state (y'', q') (where his demand is $-y''$), which would have yielded the negations of[27] (6) and (7).

Utility functions that satisfy these conditions, for $y' = 20$, $y'' = 30$, $q' = 5$, and $q'' = 9$, are

$$u^1(x_1, x_2) = \log(x_1 + a) + \log(x_2 + b),$$

$$u^2(x_1, x_2) = \log(x_1 + c) + \log(x_2 + d),$$

where $a = 0$, $b = 300$, $c \approx 32.74$, and $d \approx -65.35$, in the regions $(10, -200) \le x \le (40, -90)$ for u^1 and $(-31, 90) \le x \le (-10, 190)$ for u^2 (computations done using Maple).

The unique Walrasian equilibrium has $y^1 = -y^2 = (bc - ad)/2(b + d) \approx 10.46$ and $q = (b + d)/(a + c) \approx 7.17$. □

We point out that all the examples above are robust—small perturbations do not affect them—and have, in the relevant regions, standard Cobb–Douglas preferences that are strictly convex and smooth.

5. Socially concave games

Following the study of general economies in the previous section, we will consider the special case of economies with quasilinear utilities, where the results turn out to be quite different (see Section 7). To this end, we devote this section and the next one to the study of some relevant classes of games with convex action spaces.

This section deals with the class of "socially concave games," introduced by Even-Dar et al. (2009), who showed that in these games the time-averages of no-regret dynamics converge to Nash equilibria. The property of socially concave games that drives this result is that the expectations of correlated equilibria—and even of the more general Hannan equilibria—are pure Nash equilibria; see Theorem 9 below.

Let $\Gamma = (N, (A^i)_{i \in N}, (g^i)_{i \in N})$ be a game strategic form, where N is a finite set of *players*, and for each player $i \in N$, the set A^i of *actions* of player i is a nonempty convex subset of some Euclidean space, and $g^i : A \to \mathbb{R}$ is the payoff function of player i, where $A := \prod_{i \in N} A^i$.

The game is called a *socially concave game* (Even-Dar et al., 2009) if for each player $i \in N$ there is $\lambda^i > 0$ such that:

(G1) The function $\sum_{i \in N} \lambda^i g^i(a)$ is a concave function of $a \in A$.
(G2) For every $i \in N$ and every $a^i \in A^i$, the function $g^i(a^i, a^{-i})$ is a convex function of $a^{-i} \in A^i$.

Remarks. *(a)* (G1) and (G2) imply that $g^i(a^i, a^{-i})$ is a concave function of a^i (cf. Even-Dar et al., 2009, Lemma 2.2) for each i and a^{-i}, since $\lambda^i g^i(a^i, a^{-i}) = \sum_{j \in N} \lambda^j g^j(a^i, a^{-i}) - \sum_{j \in N \setminus i} \lambda^j g^j(a^i, a^{-i})$, and so, as a function of a^i, it is the difference between a concave function (by (G1)) and a convex function (by (G2)), and thus concave; a socially concave game is thus a so-called "concave" game (Rosen, 1965).

(b) Two payoff functions g^i and $\hat{g}^i$ are *strategically equivalent* if $\hat{g}^i(a) = \gamma^i g^i(a) + \phi^i(a^{-i})$ for all $a = (a^i, a^{-i}) \in A$, where $\gamma^i > 0$ and $\phi^i : A^{-i} \to \mathbb{R}$ is an arbitrary function that depends *only* on the actions of the *other players*. Replacing g^i with a strategically equivalent $\hat{g}^i$ does not affect the mixed best-reply correspondence and hence any strategic concepts based on best replies, such as Nash equilibrium, correlated equilibrium, and Hannan-consistent equilibrium. It also does not affect concepts based on payoff differences where the ϕ^i terms cancel, such as regrets and regret-based dynamics (see Hart and Mas-Colell, 2000, 2013, and Section 6 below). However, replacing the g^i-s with strategically equivalent $\hat{g}^i$-s may help satisfy the conditions (G1) and (G2) (as we do in Section 7).

A useful generalization of correlated equilibria is that of a *Hannan-consistent equilibrium*, or *Hannan equilibrium* for short (also known as *coarse correlated equilibrium*); see Hannan (1957), Moulin and Vial (1978), Hart and Mas-Colell (2000, 2001, 2003; 2013, Chapters 2, 3, 5), Young (2004): it is a random variable $\mathbf{a}$ with values in A, the set of action profiles, that satisfies

$$\mathbb{E}\left[g^i(\mathbf{a})\right] \ge \mathbb{E}\left[g^i(b^i, \mathbf{a}^{-i})\right] \tag{8}$$

[27] If his signal did not distinguish between (y', q') and (y', q'') then we would get the negation of (6), and if it did, then we would get both $-y' \in \arg\max_y u^2(y, -q'y)$ and $-y' \in \arg\max_y u^2(y, -q''y)$, which would imply the negation of (6).

for every player $i \in N$ and every action $b^i \in A^i$. Thus, no player i can gain by replacing his actions with any *constant* action $b^i \in A^i$, while the other players make no changes. Clearly, every correlated equilibrium is a Hannan equilibrium (take expectation over $\mathbf{s}^i$ in (2)), and the (mixed) Nash equilibria are precisely those Hannan equilibria where the actions are independent across the players, i.e., $\mathbf{a}^1, \mathbf{a}^2, \ldots, \mathbf{a}^n$ are independent random variables. In terms of distributions, let $\mu \in \Delta(A)$ be the distribution of $\mathbf{a}$; then (8) becomes

$$\int_A g^i(a)\,\mathrm{d}\mu(a) \geq \int_A g^i(b^i, a^{-i})\,\mathrm{d}\mu(a) = \int_{A^{-i}} g^i(b^i, a^{-i})\,\mathrm{d}\mu^{-i}(a^{-i}),$$

where $\mu^{-i} \in \Delta(A^{-i})$ denotes the marginal distribution of μ on A^{-i}; we refer to μ as a *Hannan distribution*.

The main property of socially concave games is the following.

Theorem 9. *Let Γ be a socially concave game. Let $\mathbf{a}$ be a Hannan equilibrium of Γ, and let $\bar{a} := \mathbb{E}[\mathbf{a}] \in A$ be its expectation. Then $\bar{a}$ is a pure Nash equilibrium of Γ, and moreover its payoffs are the same as the (expected) payoffs of $\mathbf{a}$, i.e., for every player $i \in N$ we have $g^i(\bar{a}) = \mathbb{E}\left[g^i(\mathbf{a})\right]$.*

Proof. Without loss of generality assume that the $\lambda^i = 1$ for all i. Since A is convex, $\bar{a} \in A$. Using (G2) implies

$$\mathbb{E}\left[g^i(b^i, \mathbf{a}^{-i})\right] \geq g^i\left(b^i, \mathbb{E}\left[\mathbf{a}^{-i}\right]\right) = g^i(b^i, \bar{a}^{-i}),$$

and so (8) yields

$$\mathbb{E}\left[g^i(\mathbf{a})\right] \geq g^i(b^i, \bar{a}^{-i}). \tag{9}$$

Taking in particular $b^i = \bar{a}^i$ and then summing over all players $i \in N$:

$$\sum_{i\in N} \mathbb{E}\left[g^i(\mathbf{a})\right] \geq \sum_{i\in N} g^i(\bar{a}^i, \bar{a}^{-i}) = \sum_{i\in N} g^i(\bar{a}).$$

Using (G1) implies

$$\sum_{i\in N} \mathbb{E}\left[g^i(\mathbf{a})\right] = \mathbb{E}\left[\sum_{i\in N} g^i(\mathbf{a})\right] \leq \sum_{i\in N} g^i(\mathbb{E}[\mathbf{a}]) = \sum_{i\in N} g^i(\bar{a}).$$

Therefore

$$\sum_{i\in N} \mathbb{E}\left[g^i(\mathbf{a})\right] = \sum_{i\in N} g^i(\bar{a}),$$

and so we must have equalities throughout. Thus $g^i(\bar{a}) = \mathbb{E}\left[g^i(\mathbf{a})\right]$ for every i, and (9) becomes

$$g^i(\bar{a}) \geq g^i(b^i, \bar{a}^{-i})$$

for every $b^i \in A^i$; therefore $\bar{a}$ is a pure Nash equilibrium of Γ. □

Remark. Theorem 9 gives the result of Even-Dar et al. (2009)[28]: for socially concave games, the time-average of any dynamic leading to correlated or Hannan equilibria yields a dynamic leading to pure Nash equilibria. In particular, for no-regret dynamics, the time-average of the empirical distributions converges to the set of pure Nash equilibria.

Under strict concavity we get also a uniqueness result. A socially concave game Γ will be called a socially *strictly* concave game if it satisfies (in addition to (G2)) the strict version of (G1):

(G1s) The function $\sum_{i\in N} \lambda^i g^i(a)$ is a strictly concave function of $a \in A$.

Proposition 10. *Let Γ be a socially strictly concave game. If a Hannan equilibrium exists*[29] *then it is unique, and moreover it is the unique pure Nash equilibrium (and is thus also the unique mixed Nash equilibrium and the unique correlated equilibrium).*

28 In fact, it was motivated by this result.

29 In this section we have not assumed that the action spaces A^i are compact, and so existence of equilibria is not guaranteed.

Proof. Again assume without loss of generality that $\lambda^i = 1$ for all i. Let $\mathbf{a}$ be any Hannan equilibrium, and $\bar{a}$ its expectation. Theorem 9 yields $\mathbb{E}\left[\sum_{i \in N} g^i(\mathbf{a})\right] = \sum_{i \in N} g^i(\mathbb{E}[\mathbf{a}])$; since $\sum_{i \in N} g^i$ is strictly concave, it follows that $\mathbf{a}$ must be a constant, and so $\mathbf{a} \equiv \bar{a}$ (and $\bar{a}$ is a pure Nash equilibrium). This holds for any Hannan equilibrium; but the set of Hannan equilibria is convex, and so, if it is nonempty, then it must consist of a unique point, namely, a constant $a^* \in A$. The result follows because the Hannan set contains all Nash equilibria and all correlated equilibria. □

6. Deterministic regret-matching dynamics

In this section we consider (unconditional-)regret-matching dynamics (see Hart and Mas-Colell, 2000, Theorem B; 2013, Chapter 2), and show that for a suitable class of games with convex action spaces one may replace such stochastic dynamics with deterministic ones.[30]

We consider the following additional assumptions on the strategic game $\Gamma = (N, (A^i)_{i \in N}, (g^i)_{i \in N})$:

(G3) For every $i \in N$, the action set A^i is compact.

(G4) For every $i \in N$ and $a^{-i} \in A^{-i}$, the function $g^i(a^i, a^{-i})$ is a concave function of $a^i \in A^i$.

(G5) For every $i \in N$, the function $g^i(a^i, a^{-i})$ is uniformly equicontinuous in a^i; i.e., for every $\varepsilon > 0$ there is $\delta > 0$ such that $||a^i - b^i|| < \delta$ implies $|g^i(a^i, a^{-i}) - g^i(b^i, a^{-i})| < \varepsilon$, for all $a^i, b^i \in A^i$ and $a^{-i} \in A^{-i}$.

Remark. If the function $g^i : A^i \times A^{-i} \to \mathbb{R}$ can be extended to a function $\tilde{g}^i : C^i \times A^{-i} \to \mathbb{R}$ where $C^i \supset A^i$ is an open convex set, such that $\tilde{g}^i$ is convex in $a^i \in C^i$ for every fixed $a^{-i} \in A^{-i}$ and is bounded in $a^{-i} \in A^{-i}$ for every fixed $a^i \in C^i$, then condition (G5) is satisfied, as $\tilde{g}^i$ is even uniformly equi-Lipschitz (i.e., there is L such that $|g^i(a^i, a^{-i}) - g^i(b^i, a^{-i})| \le L||a^i - b^i||$ for all $a^i, b^i \in A^i$ and $a^{-i} \in A^{-i}$); see Rockafellar (1970, Theorem 10.6).

Fix a probability measure ν^i on A^i that is equivalent to the Lebesgue measure[31] (i.e., each one is absolutely continuous with respect to the other; we could well take ν^i to be the Lebesgue measure itself on A^i, normalized so that the total mass is 1). As we will see below, the specific choice of ν^i does not matter.

Consider a repeated play of the game Γ, where $a_t = (a_t^i)_{i \in N}$ is the N-tuple of actions played at time t, for $t = 1, 2, \ldots$. For each time period $T = 1, 2, \ldots$, player $i \in N$, and pure action $b^i \in A^i$, the *unconditional regret*[32] $R_T^i(b^i)$ of b^i is given by

$$R_T^i(b^i) := \left[D_T^i(b^i)\right]_+, \quad \text{where}$$

$$D_T^i(b^i) := \frac{1}{T} \sum_{t=1}^{T} \left[g^i(b^i, a_t^{-i}) - g^i(a_t)\right].$$

We define a simple strategy[33] for player i in the repeated play of the game Γ, which we will call *DURM* (short for *Deterministic Unconditional Regret-Matching*), as follows: at period 1 or at any period $T + 1$ where $R_T^i(b^i) = 0$ for ν^i-almost every b^i, play arbitrarily; otherwise, play $a_{T+1}^i \in A^i$ given by

$$a_{T+1}^i := \frac{1}{\int_{A^i} R_T^i(b^i)\, d\nu^i(b^i)} \int_{A^i} b^i R_T^i(b^i)\, d\nu^i(b^i)\,. \tag{10}$$

Unconditional regret-matching, a stochastic strategy introduced in Hart and Mas-Colell (2000, (4.2); 2013, Chapter 2) for games with finite action spaces, plays all the pure actions b^i with probabilities that are directly proportional to their regrets, i.e., $R_T^i(b^i)/\sum_{a^i} R_T^i(a^i)$. Here, where A^i is a convex set, we replace this random play by its expectation—which yields formula (10).

Remark. The specific measure ν^i on A^i, as well as the arbitrary choices at $T = 1$ and when all regrets vanish, do affect the sequence of actions played, but do not matter in the limit, where convergence to the Hannan set obtains (see below).[34]

We now prove that playing DURM makes the regrets vanish in the limit. We start with L_2-convergence: let $\left\|R_T^i\right\|_2 = \left(\int_{A^i} \left[R_T^i(b^i)\right]^2 d\nu^i(b^i)\right)^{1/2}$ be the L_2-norm of R_T^i (cf. Hart and Mas-Colell, 2000; 2013, Chapter 2).

30 It would be interesting to study also other classes of dynamics that yield no regret (e.g., see the book of Cesa-Bianchi and Lugosi, 2006, and the gradient-ascent method of Zinkevich, 2003). Some of the relevant issues are speed of convergence, robustness, efficiency of computation, and degree of naturalness.

31 This is the d-dimensional Lebesgue measure, where d is the dimension of the affine Euclidean space generated by A^i.

32 Also known as *external* regret.

33 We use the term "action" for the one-shot game, and "strategy" for the repeated game.

34 Changing the probability weights given by ν^i is similar, for instance, to duplicating an action in the case of a finite action space.

Proposition 11. *Let the game Γ satisfy (G3) and (G4). If player $i \in N$ plays a DURM strategy then*

$$\left\| R_T^i \right\|_2 \leq \frac{2M}{\sqrt{T}} \to_{T\to\infty} 0$$

for any strategies[35] *of the other players $j \neq i$, where $M := \sup_{a\in A} |g^i(a)|$.*

Proof. We will prove this[36] for any pure strategies of $N\backslash\{i\}$, from which it immediately follows for mixed strategies too. We have

$$\begin{aligned}(T+1)D_{T+1}^i(b^i) &= TD_T^i(b^i) + \left[g^i(b^i, a_{T+1}^{-i}) - g^i(a_{T+1})\right]\\ &\leq TR_T^i(b^i) + \left[g^i(b^i, a_{T+1}^{-i}) - g^i(a_{T+1})\right].\end{aligned}$$

Squaring this inequality and noting that $\left[R_{T+1}^i(b^i)\right]^2$ equals either $\left[D_{T+1}^i(b^i)\right]^2$ or 0, and so $\left[R_{T+1}^i(b^i)\right]^2 \leq \left[D_{T+1}^i(b^i)\right]^2$, yields

$$\begin{aligned}(T+1)^2\left[R_{T+1}^i(b^i)\right]^2 &\leq T^2\left[R_T^i(b^i)\right]^2 + 4M^2\\ &\quad + 2TR_T^i(b^i)\left[g^i(b^i, a_{T+1}^{-i}) - g^i(a_{T+1})\right].\end{aligned}$$

Integrating over $b^i \in A^i$ gives

$$\begin{aligned}(T+1)^2\left\|R_{T+1}^i\right\|_2^2 &\leq T^2\left\|R_T^i\right\|_2^2 + 4M^2\\ &\quad + 2T\int_{A^i} R_T^i(b^i)\left[g^i(b^i, a_{T+1}^{-i}) - g^i(a_{T+1})\right] \mathrm{d}\nu^i(b^i).\end{aligned} \tag{11}$$

When $R_T^i(\cdot) = 0$ (a.s.) the last term vanishes; otherwise, the definition (10) of a_{T+1}^i as an average of b^i-s and the concavity of $g^i(\cdot, a_{T+1}^{-i})$ on A^i imply that

$$g^i(a_{T+1}^i, a_{T+1}^{-i}) \geq \frac{\int_{A^i} g^i(b^i, a_{T+1}^{-i}) R_T^i(b^i)\, \mathrm{d}\nu^i(b^i)}{\int_{A^i} R_T^i(b^i)\, \mathrm{d}\nu^i(b^i)},$$

and so the last term in (11) is ≤ 0. Therefore, putting $\rho_T := T^2\left\|R_T^i\right\|_2^2$, we have $\rho_{T+1} \leq \rho_T + 4M^2$, hence $\rho_T \leq 4M^2 T$, which completes the proof. □

With the additional uniform equicontinuity condition (G5), we obtain uniform convergence of the regrets.[37]

Theorem 12. *Let the game Γ satisfy (G3), (G4), and (G5). If player $i \in N$ plays a DURM strategy, then the regrets R_T^i converge uniformly to 0, i.e., $\max_{b^i\in A^i} R_T^i(b^i) \to_{T\to\infty} 0$, for any strategies of the other players $j \neq i$.*

Proof. First, note that the uniform equicontinuity of the payoff function g^i implies the equicontinuity of the regret functions: if $||b^i - c^i|| < \delta$ then

$$|R_T^i(b^i) - R_T^i(c^i)| \leq |D_T^i(b^i) - D_T^i(c^i)| \leq \frac{1}{T}\sum_{t=1}^{T} |g^i(b^i, a_t^{-i}) - g^i(c^i, a_t^{-i})| < \varepsilon.$$

Therefore the sequence R_T^i satisfies the conditions of the Arzelà–Ascoli Theorem. If R_0 is any limit point of the sequence (with respect to uniform convergence), then, by Proposition 11, $||R_0||_2 = 0$, and so $R_0 = 0$ Lebesgue-a.e., hence everywhere (since R_0 is continuous).[38] □

---

35 The strategies of the players $j \in N\backslash i$ may be independent or correlated.

36 This is essentially the proof of no regret via Blackwell (1956) approachability (see Hart and Mas-Colell, 2000, 2001; 2013, Chapters 2 and 3), but with randomizations replaced by their expectations.

37 We thank Yishay Mansour for suggesting this result.

38 A more precise analysis shows that the rate of convergence is polynomial in T (for instance, $O(T^{-1/2(d+1)})$, where d is the dimension of the affine space generated by A^i).

Given the sequence $a_t \in A$ of played action profiles, for every $T \geq 1$ let $\alpha_T \in \Delta(A)$ be the distribution that puts probability $1/T$ on each one of $a_1, a_2, \ldots, a_T$—this is the *empirical distribution of play*—and let $\bar{a}_T := (1/T)\sum_{t=1}^{T} a_t \in A$ denote its expectation—this is the *time-average play.*[39]

Corollary 13. *Let the game Γ satisfy (G3), (G4), and (G5). If each player plays a DURM strategy then the sequence α_T of empirical distributions of play converges as $T \to \infty$ to the set of Hannan distributions of Γ. Moreover, if Γ is a socially concave game (i.e., it satisfies also (G1) and (G2)), then the sequence of time-average plays $\bar{a}_T$ converges as $T \to \infty$ to the set of pure Nash equilibria of Γ.*

Proof. Let $\alpha_0 \in \Delta(A)$ be any limit point of the sequence α_T, say $\alpha_{T'} \to \alpha_0$. Therefore $D^i_{T'}(b^i) = \int_A [g^i(b^i, a^{-i}) - g^i(a)]\mathrm{d}\alpha_{T'}(a) \to \int_A [g^i(b^i, a^{-i}) - g^i(a)]\mathrm{d}\alpha_0(a)$ (since g^i is a continuous function). Theorem 12 implies that this limit, which we denote by $D^i_{\alpha_0}(b^i)$, is ≤ 0, and so $[D^i_{\alpha_0}(b^i)]_+ = 0$: the regrets at α_0 vanish for every $i \in N$ and $b^i \in A^i$. Therefore α_0 is a Hannan distribution. The "moreover" statement follows from Theorem 9. $\square$

When the game Γ has a unique Hannan equilibrium, which is therefore also the unique pure Nash equilibrium of Γ (see for instance Proposition 10), and the payoffs are strictly concave in one's own action, we get a stronger result: the period-by-period actual play $a_t \in A$ converges to the unique pure Nash equilibrium of the game. Consider the following assumption:

(G4s) For every $i \in N$ and $a^{-i} \in A^{-i}$, the function $g^i(a^i, a^{-i})$ is a strictly concave function of $a^i \in A^i$.

To get convergence of the single-period play (rather than the time-average), we need to specify the choice when all regrets vanish. Let *DURM*$_0$ be the variant of DURM with "inertia" that plays the same action as in the previous period, i.e., $a^i_{T+1} = a^i_T$, when $R^i_T(\cdot) = 0$ (a.s.).

Proposition 14. *Let the game Γ be a socially strictly concave game (i.e., it satisfies (G1s) and (G2)) that also satisfies (G3), (G4s), and (G5). If each player plays the DURM$_0$ strategy then the sequence of plays a_t converges as $t \to \infty$ to the unique pure Nash equilibrium of Γ.*

Proof. Let $\hat{a} \in A$ be the unique pure Nash equilibrium, which is also the unique Hannan equilibrium (see Proposition 10). Given $\varepsilon > 0$, let $\delta > 0$ be such that all the δ-best replies to $\hat{a}^{-i}$ are within ε of $\hat{a}^i$ (such a δ exists since the best reply is unique—by (G4s)—and the A^i are compact (G3)). Next, $D^i_T(b^i) = \int_A [g^i(b^i, a^{-i}) - g^i(a)]\mathrm{d}\alpha_T(a) \to g^i(b^i, \hat{a}^{-i}) - g^i(\hat{a})$ (since the distribution α_T converges to the Dirac measure on $\hat{a}$ by Corollary 13, and g^i is continuous); moreover, this convergence is uniform in b^i (by (G5)), and so there is T_0 such that $D^i_T(b^i) < [g^i(b^i, \hat{a}^{-i}) - g^i(\hat{a})] + \delta$ for all $T \geq T_0$, all $i \in N$, and all $b^i \in A^i$. Therefore $D^i_T(b^i) > 0$ implies that $g^i(b^i, \hat{a}^{-i}) - g^i(\hat{a}) \geq -\delta$, and so b^i is a δ-best reply to $\hat{a}^{-i}$; hence $||b^i - \hat{a}^i|| \leq \varepsilon$. If $R^i_T(\cdot)$ does not vanish (a.s.), then the action a^i_{T+1} that DURM plays in this case is an average of b^i-s with $D^i_T(b^i) > 0$ (see (10)), and so it follows that $||a^i_{T+1} - \hat{a}^i|| \leq \varepsilon$. Together with the "inertia" condition of DURM$_0$, if $R^i_{T_1}(\cdot)$ does not vanish for some $T_1 \geq T_0$ then we have $||a^i_{T+1} - \hat{a}^i|| < \varepsilon$ for all $T \geq T_1$. Otherwise $R^i_T(\cdot) = 0$ (a.s.) for all $T \geq T_0$, and then the play is constant: $a^i_{T+1} = a^i_{T_0}$; but the empirical distribution α_T converges to $\mathbf{1}_{\hat{a}}$, and so we must have $a^i_{T_0} = \hat{a}^i$, and the proof is complete. $\square$

Remarks. *(a)* The proof above shows that the result holds for any strategies that guarantee uniform convergence of the regrets to 0 and that plays only actions that have positive regret when such actions exist.

(b) The results of this section can be extended to the larger class of *regret-based* strategies (cf. Hart and Mas-Colell, 2001; 2013, Chapter 3).

6.1. Deterministic conditional-regret dynamics: an impossibility result

The existence of a *deterministic* strategy that guarantees that the unconditional regret vanishes in the limit may come as a surprise, particularly since this is not the case for conditional regret (recall that having no conditional regret corresponds to correlated equilibria; see Hart and Mas-Colell, 2000; 2013, Chapter 2). Indeed, let

$$\hat{R}^i_T(a^i \to b^i) := \left[\frac{1}{T} \sum_{t \leq T : a^i_t = a^i} \left[g^i(b^i, a^{-i}_t) - g^i(a_t) \right] \right]_+ ,$$

where[40] $a^i, b^i \in A^i$, be the *conditional regret* of player i from a^i to b^i; see Hart and Mas-Colell (2000; 2013, Chapter 2).

[39] Thus $\alpha_T \in \Delta(A)$ and $\bar{a}_T \in A$.

[40] If an action a^i has not been played up to time T then $\hat{R}^i_T(a^i \to b^i) = 0$ for every b^i (as the sum is empty).

Proposition 15. *Assume that player $i \in N$ does not have a weakly dominant action in the game Γ (i.e., there is no $a_0^i \in A^i$ such that $g^i(a_0^i, a^{-i}) \geq g^i(a^i, a^{-i})$ for every $a^i \in A^i$ and $a^{-i} \in A^{-i}$) and that g^i is a continuous function on A. Then for every deterministic strategy of i in the repeated play of Γ, the conditional regrets of i* cannot *be guaranteed to converge to* 0*. More precisely: there exists a $\delta > 0$ and a (deterministic) strategy of the other players $N \backslash \{i\}$ such that*[41]

$$\sum_{a^i \in A^i} \max_{b^i \in A^i} \hat{R}_T^i(a^i \to b^i) \geq \delta$$

for all $T \geq 1$.

Proof. For every $a^i \in A^i$ let $F^{-i}(a^i) \in A^{-i}$ be such that

$$\max_{b^i \in A^i} g^i(b^i, F^{-i}(a^i)) > g^i(a^i, F^{-i}(a^i))$$

(if there were no such $F^{-i}(a^i)$ then a^i would be a dominant action of i); the continuity of g^i and the compactness of A^i and A^{-i} imply that there is a $\delta > 0$ such that the difference in the above inequality is at least δ for all $a^i \in A^i$, i.e.,

$$\max_{b^i \in A^i} g^i(b^i, F^{-i}(a^i)) - g^i(a^i, F^{-i}(a^i)) \geq \delta.$$

Let $H^i(a^i) \in A^i$ be an action of i satisfying this inequality; i.e., for every $a^i \in A^i$ we have

$$g^i(H^i(a^i), F^{-i}(a^i)) - g^i(a^i, F^{-i}(a^i)) \geq \delta. \tag{12}$$

Take a deterministic strategy of player i. Let $a^i \in A^i$ be the pure action played by player i at time t after some history h_{t-1}, and let the other players play $F^{-i}(a^i)$ after that same history h_{t-1}. Since every time that i plays a^i the other players play $F^{-i}(a^i)$, for every $T \geq t$ we then get

$$\hat{R}_T^i(a^i \to b^i) = \phi_T^i(a^i)\left[g^i(b^i, F^{-i}(a^i)) - g^i(a^i, F^{-i}(a^i))\right]_+,$$

where $\phi_T^i(a^i)$ denotes the relative frequency that a^i has been played up to time T, and so

$$\max_{b^i \in A^i} \hat{R}_T^i(a^i \to b^i) \geq \phi_T^i(a^i)\delta.$$

Since $\sum_{a^i} \phi_T^i(a^i) = 1$ we get $\sum_{a^i \in A^i} \max_{b^i \in A^i} \hat{R}_T^i(a^i \to b^i) \geq \delta$. $\square$

Remark. When conditional regrets converge to 0, the empirical distribution of play $\alpha_T \in \Delta(A)$, which puts weight $1/T$ on each one of $a_1, a_2, \ldots, a_T \in A$, converges to the set of correlated equilibrium distributions (Hart and Mas-Colell, 2000; 2013, Chapter 2). By contrast, the strategies above keep the empirical distributions far from being correlated equilibrium distributions, for *all* T; indeed, having player i replace each "recommendation" a^i by $H^i(a^i) \in A^i$ (see (12)) increases his payoff by at least δ.

7. Economies with quasilinear utilities

In this section we consider the special class of economies *with quasilinear utilities*, i.e., where the utilities of all agents are linear in the "numeraire" good (also called *economies with transferable utility*, as the numeraire can serve to "transfer" utility between the agents.)[42] Formally, in a *quasilinear economy*, for every agent $i \in N$ there is function $v^i : \mathbb{R}^{M-1} \to \mathbb{R}$ such that $u^i(x_1, \ldots, x_{m-1}, x_m) = v^i(x_1, \ldots, x_{m-1}) + x_m \in \mathbb{R}^{M-1}$ for every[43] $x = (x_1, \ldots, x_m) \in \mathbb{R}^M$. The general definition of a Walrasian equilibrium becomes: a price vector $\tilde{q} \in \mathbb{R}_{++}^{M-1}$ for the goods $M \backslash \{m\}$ (recall that the price of good m is normalized to 1) and an allocation $(\tilde{y}^i)_{i \in N}$ with $\tilde{y}^i \in \mathbb{R}^{M-1}$ and $\sum_{i \in N} \tilde{y}^i = 0$ such that $\tilde{y}^i$ is a demand of i at $\tilde{q}$, for every $i \in N$, i.e., $\tilde{y}^i \in \arg\max_{y^i \in \mathbb{R}^{M-1}} \left(v^i(y) - \tilde{q} \cdot y\right)$.

In the quasilinear case, the payoff functions of the associated game $\mathcal{G}$ of Section 3 (with the numeraire m as the special good) become

$$g^i(y^1, \ldots, y^{n-1}, q) := v^i(y^i) - q \cdot y^i \quad \text{for } i \neq n,$$
$$g^n(y^1, \ldots, y^{n-1}, q) := v^n(y^n) - q \cdot y^n - \sup_{w \in \mathbb{R}^{M-1}} \{v^n(w) - q \cdot w\}.$$

---

41 The sum is over those $a^i \in A^i$ that have been played up to time T, since for any other a^i the regrets are 0.

42 See Mas-Colell et al. (1995); quasilinearity is a standard hypothesis, particularly in the extensive mechanism design literature.

43 We maintain the basic assumptions of Section 2, and so v^i is concave and, for every supporting price vector $q \in \mathbb{R}^{M-1}$, we have $1/K \leq q_\ell \leq K$ and $1/K \leq q_\ell / q_{\ell'} \leq K$ for all $1 \leq \ell, \ell' \leq m-1$.

For our specific game $\mathcal{G}$, a correlated equilibrium is a random variable $(\mathbf{y}^1, \ldots, \mathbf{y}^{n-1}, \mathbf{q})$ with values in $\mathbb{R}^{M-1} \times \ldots \times \mathbb{R}^{M-1} \times Q$. Each realization generates an outcome (by (1)), and so we have the induced random outcome $(\mathbf{x}^1, \ldots, \mathbf{x}^n; \mathbf{p})$, which yields the *(expected) outcome* $(\mathbb{E}\left[\mathbf{x}^1\right], \ldots, \mathbb{E}[\mathbf{x}^n]; \mathbb{E}[\mathbf{p}])$. Let CEO be the set of *Correlated Equilibrium (expected) Outcomes* of the game $\mathcal{G}$.

It turns out that in the quasilinear case the result of Theorem 1 on the equivalence between the Walrasian and the pure Nash equilibrium outcomes becomes much stronger.

Theorem 16. *Let $\mathcal{E}$ be a quasilinear economy. Then the set of correlated equilibrium outcomes, the set of mixed Nash equilibrium outcomes, the set of pure Nash equilibrium outcomes, and the set of Walrasian outcomes, all coincide:* CEO = NEO = PNEO = WEO.

We first show that quasilinear economies yield socially concave games. Theorem 9 then implies that the expectation of a correlated equilibrium is a pure Nash equilibrium. However, that is not enough: the payoff functions are not linear, and so we need to show also that the expected outcome is the same as the outcome of the expected action profile.

Proposition 17. *The strategic game $\mathcal{G}$ generated from a quasilinear economy $\mathcal{E}$ is strategically equivalent to a socially concave game.*

Proof. Recalling Remark (b) that follows (G1)–(G2) in Section 5, take $\phi^n(y^1, \ldots, y^{n-1}) = -v^n(-\sum_{i \neq n} y^i) = -v^n(y^n)$ (and $\phi^i = 0$ for $i \neq n$). The resulting functions $\hat{g}^i$ are convex in a^{-i} (in fact, linear), and $\sum_i \hat{g}^i = \sum_{i \neq n} v^i(y^i) + \inf_w \{q \cdot w - v^n(w)\}$ is concave in $(y^1, \ldots, y^{n-1}, q)$. □

Proof of Theorem 16. Let $\mathbf{a} = (\mathbf{y}^1, \ldots, \mathbf{y}^{n-1}, \mathbf{q})$ be a correlated equilibrium of $\mathcal{G}$, with induced random outcome $\theta = (\mathbf{x}^1, \ldots, \mathbf{x}^n, \mathbf{p})$. Let $\bar{a} = (\bar{y}^1, \ldots, \bar{y}^{n-1}, \bar{q}) := \mathbb{E}[\mathbf{a}]$ and $\bar{\theta} = (\bar{x}^1, \ldots, \bar{x}^n, \bar{p}) := \mathbb{E}[\theta]$ be the expected action profile and the expected outcome, respectively. We have to show that $\bar{\theta}$ is a pure Nash equilibrium outcome. Proposition 17 implies that we can apply Theorem 9, and thus $\bar{a}$ is a pure Nash equilibrium of $\mathcal{G}$, and

$$g^i(\bar{a}) = \mathbb{E}\left[g^i(\mathbf{a})\right] \tag{13}$$

for all $i \in N$. We will prove that the outcome induced by the expected action $\bar{a}$ is precisely the expected outcome $\bar{\theta}$, and so indeed $\bar{\theta} \in$ PNEO. The issue is the nonlinearity of the terms $q \cdot y^i$: we have $\bar{x}^i = \mathbb{E}\left[\mathbf{x}^i\right] = \mathbb{E}\left[(\mathbf{y}^i, -\mathbf{q} \cdot \mathbf{y}^i)\right] = \left(\bar{y}^i, -\mathbb{E}\left[\mathbf{q} \cdot \mathbf{y}^i\right]\right)$, whereas the consumption of agent i in the outcome induced by $\bar{a}$ is $(\bar{y}^i, -\bar{q} \cdot \bar{y}^i)$, and so we have to show that $\mathbb{E}\left[\mathbf{q} \cdot \mathbf{y}^i\right] = \bar{q} \cdot \bar{y}^i$ for all i.

For agents $i \neq n$, (13) is $\mathbb{E}\left[v^i(\mathbf{y}^i) - \mathbf{q} \cdot \mathbf{y}^i\right] = v^i(\bar{y}^i) - \bar{q} \cdot \bar{y}^i$. Since v^i is concave we have $\mathbb{E}\left[v^i(\mathbf{y}^i)\right] \leq v^i\left(\mathbb{E}\left[\mathbf{y}^i\right]\right) = v^i(\bar{y}^i)$, and so

$$\mathbb{E}\left[\mathbf{q} \cdot \mathbf{y}^i\right] \leq \bar{q} \cdot \bar{y}^i. \tag{14}$$

For agent n, we recall Lemma 2. First, we have $g^n(\bar{a}) = 0$ (since $\bar{a}$ is a Nash equilibrium); and second, we get $g^n(\mathbf{a}) = 0$ (a.s.) (since $g^n \leq 0$ and $0 = g^n(\bar{a}) = \mathbb{E}[g^n(\mathbf{a})]$ by (13)). Thus $\mathbf{y}^n$ is a demand of agent n at $\mathbf{q}$ (a.s.), and so, in particular, it is weakly preferred to $\bar{y}^n = \mathbb{E}[\mathbf{y}^n]$, i.e., $v^n(\bar{y}^n) - \mathbf{q} \cdot \bar{y}^n \leq v^n(\mathbf{y}^n) - \mathbf{q} \cdot \mathbf{y}^n$. Taking expectation yields $v^n(\bar{y}^n) - \bar{q} \cdot \bar{y}^n \leq \mathbb{E}[v^n(\mathbf{y}^n) - \mathbf{q} \cdot \mathbf{y}^n]$. The right-hand side is $\leq v^n(\bar{y}^n) - \mathbb{E}[\mathbf{q} \cdot \mathbf{y}^n]$ by the concavity of v^n, and so we obtain the inequality (14) also for $i = n$.

Summing (14) over all $i \in N$ yields $\mathbb{E}\left[\mathbf{q} \cdot \sum_i \mathbf{y}^i\right] \leq \bar{q} \cdot \sum_i \bar{y}^i = \bar{q} \cdot \mathbb{E}\left[\sum_i \mathbf{y}^i\right]$, which is in fact an equality, as both sides equal 0 (because $\sum_i \mathbf{y}^i = 0$). This implies that all the inequalities above, in particular (14), become equalities, and so CEO = PNEO. Together with CEO ⊃ NEO ⊃ PNEO, and then Theorem 1 that gives PNEO = WEO, the proof is complete. □

A similar result holds for sunspot equilibria. The *(expected) outcome* of a sunspot equilibrium is the expectation of the random outcome $(\mathbf{x}^1, \ldots, \mathbf{x}^n; \mathbf{p})$; i.e., it is the outcome $(\mathbb{E}\left[\mathbf{x}^1\right], \ldots, \mathbb{E}[\mathbf{x}^n]; \mathbb{E}[\mathbf{p}])$. Let SEO denote the set of *Sunspot Equilibrium (expected) Outcomes*.

Proposition 18. *Let $\mathcal{E}$ be a quasilinear economy. Then the set of sunspot equilibrium outcomes coincides with the set of Walrasian equilibrium outcomes:* SEO = WEO.

Proof. Consider a sunspot equilibrium as in Section 4.2 and an agent $i \in N$. With $\bar{y}^i := \mathbb{E}\left[\mathbf{y}^i\right]$ and $\bar{q} := \mathbb{E}[\mathbf{q}]$, the sunspot equilibrium condition (3) yields, in particular, $\mathbb{E}\left[u^i(\mathbf{y}^i, -\mathbf{q} \cdot \mathbf{y}^i) \mid \mathbf{s}^i\right] \geq \mathbb{E}\left[u^i(\bar{y}^i, -\mathbf{q} \cdot \bar{y}^i) \mid \mathbf{s}^i\right]$, or

$$\mathbb{E}\left[v^i(\mathbf{y}^i) - \mathbf{q} \cdot \mathbf{y}^i \mid \mathbf{s}^i\right] \geq \mathbb{E}\left[v^i(\bar{y}^i) - \mathbf{q} \cdot \bar{y}^i \mid \mathbf{s}^i\right]. \tag{15}$$

Taking expectation over $\mathbf{s}^i$ gives

$$\mathbb{E}\left[v^i(\mathbf{y}^i)\right] - \mathbb{E}\left[\mathbf{q} \cdot \mathbf{y}^i\right] \geq v^i(\bar{y}^i) - \bar{q} \cdot \bar{y}^i.$$

Since v^i is concave we have $\mathbb{E}\left[v^i(\mathbf{y}^i)\right] \le v^i\left(\mathbb{E}\left[\mathbf{y}^i\right]\right) = v^i(\bar{y}^i)$, and so we get

$$\mathbb{E}\left[\mathbf{q}\cdot\mathbf{y}^i\right] \le \bar{q}\cdot\bar{y}^i. \tag{16}$$

As in the proof of Theorem 16 (sum over i and recall that $\sum_i \mathbf{y}^i = 0$), we must have equalities all along. In particular (15) is an equality, and so $\bar{y}^i$ is also a maximizer in the sunspot equilibrium condition (3). After taking expectation over $\mathbf{q}$, (3) becomes $v^i(\bar{y}^i) - \bar{q}\cdot\bar{y}^i \ge v^i(y) - \bar{q}\cdot y$ for all y, and so the outcome corresponding to $(\bar{y}^1, \dots, \bar{y}^{n-1}, \bar{q})$ is indeed a Walrasian equilibrium outcome. □

What the proofs above suggest is that while correlated equilibria and sunspot equilibria may entail some randomness, it is mostly inessential. Under some standard additional assumptions of strict concavity and smoothness, this randomness is completely eliminated and the equilibrium concepts reduce to the unique Walrasian equilibrium.

Proposition 19. *Let $\mathcal{E}$ be a quasilinear economy, and assume that all the functions v^i are strictly concave and differentiable.*[44] *If a Walrasian equilibrium exists*[45] *then it is unique, and in this case the associated game $\mathcal{G}$ has a unique pure Nash equilibrium, which is also the unique correlated equilibrium. Moreover, all sunspot equilibria yield in every state the same allocation of the non-numeraire goods and the same conditional expected prices as in the unique Walrasian equilibrium.*

Proof. We start with correlated equilibria. In the proof of Theorem 16 we have seen that all inequalities are in fact equalities, and so in particular $\mathbb{E}\left[v^i(\mathbf{y}^i)\right] = v^i\left(\bar{y}^i\right)$ (with $\bar{y}^i = \mathbb{E}\left[\mathbf{y}^i\right]$); the strict concavity of v^i implies that $\mathbf{y}^i$ must be constant, i.e., $\mathbf{y}^i = \bar{y}^i$. Then $\mathbf{q}$ is a supergradient of v^n at $\bar{y}^n$, and so the smoothness of v^n implies that $\mathbf{q}$ is constant, i.e., $\mathbf{q} = \bar{q}$. Thus there is no randomness, which means that every correlated equilibrium is a pure Nash equilibrium. Hence there is at most one pure Nash equilibrium (otherwise we would have random mixtures of different pure Nash equilibria as correlated equilibria). By Theorem 1, this implies also the uniqueness of the Walrasian equilibrium.

For sunspot equilibria, recall the proof of Proposition 18. Again, the equalities we have obtained include $\mathbb{E}\left[v^i(\mathbf{y}^i)\right] = v^i\left(\bar{y}^i\right)$, and so the strict concavity of v^i implies $\mathbf{y}^i = \bar{y}^i$. Next, the sunspot equilibrium condition (3) says that $\mathbb{E}\left[\mathbf{q}|\mathbf{s}^i\right]$ is a supergradient of v^i at $\bar{y}^i$, which is unique because v^i is a smooth function. Therefore $\mathbb{E}\left[\mathbf{q}|\mathbf{s}^i\right] = \bar{q}$. □

Remarks. *(a)* An example of a nontrivial sunspot equilibrium in the setup of Proposition 19, where $(\bar{y}^1, \dots, \bar{y}^n, \bar{q})$ is the unique Walrasian equilibrium, is as follows. Let $\mathbf{y}^i = \bar{y}^i$ for all i; let $\mathbf{q}$ take the values q' and q'' with equal probabilities, where $q' \ne q''$ satisfy $\bar{q} = (1/2)q' + (1/2)q''$; and assume no agent gets any information.

(b) In Proposition 19, if there is no Walrasian equilibrium then the proof implies that there are no Nash equilibria, no correlated equilibria, and no sunspot equilibria.

(c) All the results in this section apply not only to correlated equilibria, but also to the more general concept of Hannan-consistent equilibria; indeed, they are all based on Theorem 9 (see Section 5).

We come now to dynamics, where we will use the theory developed in Section 6. We thus need to work with compact action spaces, and so we assume that there is a finite bound b_0 on the individually rational outcomes. This bound could be given by the cone condition (C) in Section 3.1 (see Lemma 3), or by the total initial endowment of the non-numeraire goods (by carrying out a construction as in Section 3.2, but restricted to the non-numeraire goods).

Proposition 20. *Let $\mathcal{E}$ be a quasilinear economy with a bound b_0 on the individually rational outcomes. If every agent plays a DURM strategy in the game $\mathcal{G}_b$ with $b \ge b_0$, then the time-average of the action profiles converges to the set of pure Nash equilibria of the game, whose outcomes coincide with the set of Walrasian equilibrium outcomes.*

Proof. Since $\mathcal{G}_b$ satisfies (G1)–(G5) (for (G5), see the remark immediately following it in Section 6), the result follows from Corollary 13. □

In the strict case, the actual play also converges to the Walrasian outcome.

Proposition 21. *Let $\mathcal{E}$ be a quasilinear economy with a bound b_0 on the individually rational outcomes, and assume that all the functions v^i are strictly concave and differentiable. If every agent plays a DURM strategy in the game $\mathcal{G}_b$ with $b \ge b_0$, then the outcome θ_t obtained at time t converges to the unique Walrasian equilibrium outcome.*

Proof. In this case the strict assumptions (G1s) and (G4s) are also satisfied, and we apply Proposition 14. □

44 Weaker conditions may suffice: $v^1, \dots, v^{n-1}$ strictly concave and v^n differentiable.

45 We do not assert here the existence of equilibria; to guarantee this one would need some boundedness assumptions, e.g., condition (C) in Section 3.1.

Remarks. *(a)* DURM may be replaced by other dynamics that yield no regret in the limit.

(b) Assumption (C) justifies the restriction of our dynamics to a compact set (see Section 3.1; also, Section 3.2). It remains an open question whether, under condition (C), unrestricted DURM (where the regrets of all actions—not only those bounded by b—are considered) also yields the same results.

(c) It is well known that, informally speaking, exchange economies are well behaved for classes of economies other than the quasilinear class—in particular for the gross-substitute class (that is, with a low degree of complementarity), which includes the Cobb–Douglas utilities. It would certainly be interesting to study the regret-based dynamics in these classes.

References

Arrow, K.J., Debreu, G., 1954. The existence of an equilibrium for a competitive economy. Econometrica 22, 265–290.
Aumann, R.J., 1974. Subjectivity and correlation in randomized strategies. J. Math. Econ. 1, 67–96.
Blackwell, D., 1956. An analog of the minmax theorem for vector payoffs. Pac. J. Math. 6, 1–8.
Cass, D., Shell, K., 1991. Do sunspots matter? J. Polit. Economy 91, 193–227.
Cesa-Bianchi, N., Lugosi, G., 2006. Prediction, Learning, and Games. Cambridge University Press.
Debreu, G., 1959. Theory of Value. Wiley.
Even-Dar, E., Mansour, Y., Nadav, U., 2009. On the convergence of regret minimization dynamics in concave games. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing. STOC 2009, pp. 523–532.
Forges, F., Peck, J., 1995. Correlated equilibrium and sunspot equilibrium. Econ. Theory 5, 33–50.
Groves, T., Ledyard, J., 1987. Incentive compatibility since 1972. In: Groves, T., Radner, R., Reiter, S. (Eds.), Information, Incentives and Economic Mechanisms. University of Minnesota Press, pp. 48–111.
Hannan, J., 1957. Approximation to Bayes risk in repeated play. In: Dresher, M., Tucker, A.W., Wolfe, P. (Eds.), Contributions to the Theory of Games, vol. III. In: Annals of Mathematics Studies, vol. 39. Princeton University Press, pp. 97–139.
Hart, O., 1974. On the existence of equilibrium in a securities model. J. Econ. Theory 9, 293–311.
Hart, S., Mas-Colell, A., 2000. A simple adaptive procedure leading to correlated equilibrium. Econometrica 68, 1127–1150 [also: Chapter 2 in Hart and Mas-Colell (2013)].
Hart, S., Mas-Colell, A., 2001. A general class of adaptive strategies. J. Econ. Theory 98, 26–54 [also: Chapter 3 in Hart and Mas-Colell (2013)].
Hart, S., Mas-Colell, A., 2003. Regret-based continuous-time dynamics. Games Econ. Behav. 45, 375–394 [also: Chapter 5 in Hart and Mas-Colell (2013)].
Hart, S., Mas-Colell, A., 2013. Simple Adaptive Strategies. World Scientific Publishing.
Hurwicz, L., 1979. Outcome functions yielding Walrasian and Lindahl allocations at Nash equilibrium points. Rev. Econ. Stud. 46, 217–225.
Hurwicz, L., Maskin, E., Postlewaite, A., 1995. Feasible Nash implementation of social choice rules when the designer does not know endowments or production sets. In: Ledyard, J. (Ed.), The Economics of Informational Decentralization: Complexity, Efficiency, Stability. Kluwer, pp. 367–433.
Mas-Colell, A., Whinston, M., Green, J., 1995. Microeconomic Theory. Oxford University Press.
Maskin, E., Tirole, J., 1987. Correlated equilibria and sunspots. J. Econ. Theory 43, 364–373.
Moulin, H., Vial, J.P., 1978. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. Int. J. Game Theory 7, 201–221.
Polemarchakis, H., Ray, I., 2006. Sunspots, correlation, and competition. Games Econ. Behav. 56, 174–184.
Postlewaite, A., Wettstein, D., 1983. Continuous and feasible implementation. Rev. Econ. Stud. 56, 603–611.
Rockafellar, R.T., 1966. Level sets and continuity of conjugate convex functions. Trans. Amer. Math. Soc. 123, 46–63.
Rockafellar, R.T., 1970. Convex Analysis. Princeton University Press.
Rosen, J.B., 1965. Existence and uniqueness of equilibrium points for concave N-person games. Econometrica 33, 520–534.
Schmeidler, D., 1980. Walrasian analysis via strategic outcome functions. Econometrica 48, 1585–1594.
Shell, K., 2007. Sunspot equilibrium. In: Eatwell, J., Milgate, M., Newman, P. (Eds.), The New Palgrave: A Dictionary of Economics. Macmillan, pp. 274–280.
Young, H.P., 2004. Strategic Learning and Its Limits. Oxford University Press.
Zinkevich, M., 2003. Online convex programming and generalized infinitesimal gradient ascent. In: Proceedings of the Twentieth International Conference on Machine Learning. ICML-2003, pp. 928–936.