WORKING PAPERS IN ECONOMICS

No. 06/07

S. D. FLÅM, AND O. GODAL

MARKET CLEARING AND PRICE FORMATION

Market Clearing and Price Formation

S. D. Flåm* and O. Godal

April 16, 2007

ABSTRACT. Considered here is decentralized exchange of privately owned commodity bundles. Voluntary transactions take the form of repeated bilateral barters. Under broad and reasonable hypotheses the resulting process converges to competitive equilibrium. Price-taking behavior is not assumed. Prices emerge over time; they need neither be anticipated nor known at any interim stage.

Key words: exchange economies, price equilibrium, stochastic approximation.
JEL Classification: C61, C72, C90, D51, D80.

1. INTRODUCTION

Every form of real, economic behavior is shaped by trial and error. It appears desirable therefore, that theories of markets and prices mirror such features. Theory has however, found it difficult to explain how markets function and prices emerge.

Many difficulties stem from considering only perfect agents and equilibrium outcomes.[1] That perspective is somewhat exclusive, and it inverts the natural order of things. To wit, most human-like agents appear competent only *after* interim "mistakes." And trivially, markets clear only *after* trade. So, static views on market mechanisms and price formation should yield to explicit descriptions of dynamics. Whatever be the nature of such descriptions, they had better reflect that

• information is asymmetrical, and traders may prefer to conceal own evaluations;

• nonetheless, differential information often diffuses (Grossman and Stiglitz, 1976);

• many a market is largely affected by exogenous random factors, typically having non-identified probability distributions;

• much trade happens out of equilibrium, often without middle men or brokers;

• most parties hesitate in transacting large quantities at transient price levels;

• in mature markets participants need know little to take "reasonable" actions (Hayek, 1945);

• in fact, markets may partly substitute for individual rationality (Gode and Sunder, 1993);

*Corresponding author; both are the Economics Department, Bergen University, Norway. Address: Economics Department, Postboks 7802, N-5020 Bergen, Norway. Telephone: +47 55589200. Fax: +47 55589210. E-mail: {*sjur.flaam*, *odd.godal*}@econ.uib.no. We appreciate funding from the RENERGI program of the NFR.

[1]Other difficulties come with introducing a fictitious auctioneer engaged in Walrasian tâtonnement (Arrow and Hahn, 1971; Bala and Majumdar, 1992; Saari, 1985, 1995).

- equilibrium, if any, can embody bid-ask spreads; and finally,
- as few as two agents might sometimes be around.

Studied below is an archetypical but rather general market process that incorporates all these features. Moreover, it tolerates that agents aren't perfectly competent, fully foresighted, or marvelously rational.

More modestly, it presumes that each agent steadily seeks to improve his economic welfare. Instrumental for own improvement is that pairs of agents meet time and again. During any encounter the two parties explore whether some exchange of endowments could please both. Whenever they think so, suitably small commodity bundles are transferred between them. Granted decreasing utility margins, we shall identify reasonable conditions that suffice for markets to clear - and prices to emerge - ultimately. Retrospectively, the resulting equilibrium can be seen as a price-supported core solution of a cooperative game.

Our approach relates closely to behavioral and experimental economics (Smith, 1962, 1976, 1982). That is, we fully control the *economic environment* in admitting finitely many agents, specified by their endowments and preferences. The dominating *economic institution* amounts to iterated, bilateral exchange between two parties at a time. Instead of letting human players set the process in real motion, it can be simulated as an agent-based computer model (Tesfatsion and Judd, 2006). It turns out that theoretical predictions fit nicely to experimental and computational results (Klaassen et al., 2005).

We shall not model how agents communicate or negotiate. No language or list of bid-ask messages is prescribed here. Instead, we posit that every agent invariably attempts to better himself along his actual *gradient direction*. As is well known, that direction offers him steepest welfare ascent. Presumably any two agents, each knowing merely his own gradient direction, can negotiate a small and worthwhile exchange. The resulting transactions are fairly easy to understand and implement. Also, on a more technical note, they can cope with constraints, inexact computation, non-smooth objectives, and uncertainty (Ermoliev and Wets, 1998).

All this speaks for gradient-like procedures. Added to their merits comes an extra bonus, namely: they illuminate some central issues in cognitive sciences. Specifically: how do real agents view their decision problems? How is information processed? What sort of behavior facilitates and reflects individual improvement?

Of course, only experimental evidence can elucidate these questions. Casual observation indicates, however, that typical agents form local approximations and views. Moreover, they adapt to circumstances. But! these are precisely the features that characterize gradient methods (Ermoliev and Wets, 1998). In fact, their stochastic versions mirror *four* common aspects of human behavior: *First*, probability distributions are frequently hard to learn or keep in mind. *Second*, mean values (i.e., mathematical expectations) are often impossible or costly to compute. *Third*, information concerning levels and rates is readily available only at the current point. *Fourth*, while away from a steady state, most economic agents tolerate some approximations - or some cutting of corners.

The upshot is that many arguments - including several crucial ones that come from cognitive sciences - they speak in favor of gradient methods. Here, what results is a procedure that requires little competence or experience on the part of decision makers. In fact, it obliges them neither to predict nor to optimize. Broadly, nobody must ever form, use or revise any strategy.

To explain why it's expedient to separate diverse issues. Section 2 provides preliminaries, and Section 3 considers equilibrium, also dealt with in Appendix 1. Section 4 does some groundwork, pinning down that stability will prevail. Section 5, the heart of the paper, spells out the exchange process and proves convergence to equilibrium, using stochastic approximation as main vehicle (Benaim, 1996; Benveniste et al., 1990). For simplicity, interpretation is mostly coached in terms of producers or agents having quasi-linear utility. Section 6 gets around that limitation by accommodating agents with more general preferences. Section 7 illustrates the model with emissions trading under the Kyoto Protocol. Section 8 concludes with some bibliographic remarks. Proofs are found in Appendix 2. To emphasize chief structures some signposts indicate where the reader may skip forward.

2. Preliminaries

This section prepares the ground. It introduces *agents, endowments, objectives, constraint sets, constraint qualifications*, and *uncertainty* - in that order.

Along the way it also discusses the handling of non-smooth data. The reader may skip (or postpone) that material, replacing generalized derivatives with ordinary ones. If just perusing the paper, he may proceed directly to Section 3, recollecting there that agent $i \in I$ exchanges own *endowment* $e_i \in \mathbb{R}^G$ for some $x_i \in X_i \subseteq \mathbb{R}^G$ so as to maximize *concave expected payoff* $\pi_i(x_i) = E\Pi_i(x_i, \psi_i)$.

• *Agents* constitute a finite, fixed set I. Until further notice regard them as *producers* or more generally, as agents who enjoy *quasi-linear utility*.

• The *endowment* e_i of agent $i \in I$ is codified as a point in a real finite-dimensional[2] vector space $\mathbb{X}$. That space, common to all parties, is equipped with inner product $\langle \cdot, \cdot \rangle$ and associated norm $\|\cdot\|$. When diverse commodities require separate mention, let G denote the group of relevant economic goods. Any commodity vector then comes as a G-tuple $(x_g) \in \mathbb{X} = \mathbb{R}^G$.

• The *objective* of agent i is to improve own payoff or profit $\pi_i(x_i) \in \mathbb{R}$, using factor bundle $x_i \in X_i \subseteq \mathbb{X}$. Assume decreasing returns to scale, meaning that $\pi_i(\cdot)$ is *concave*.

For analytical convenience one may take $\pi_i(\cdot)$ differentiable. It shouldn't be ignored though, that presence of underlying choices, technologies or tariffs tend to make

---

[2]With a view towards potential applications, notably in finance, all arguments will fit instances where $\mathbb{X}$ is Hilbert.

π_i non-smooth. We only require therefore, that $\pi_i : \mathbb{X} \to \mathbb{R}$ *be Lipschitz continuous* with some modulus $l_i > 0$.[3]

To deal with such criteria recall that $\chi^* \in \mathbb{X}$ is declared a *supergradient* at $\chi \in \mathbb{X}$ of a proper function $f : \mathbb{X} \to \mathbb{R} \cup \{-\infty\}$, and we write $\chi^* \in \partial f(\chi)$, iff $f(\cdot) \leq f(\chi) + \langle \chi^*, \cdot - \chi \rangle$; see Rockafellar (1970). The *superdifferential* $\partial f(\chi)$ reduces to a singleton, called the *gradient*, iff f is Gâteaux differentiable at χ. A concave function $f : \mathbb{X} \to \mathbb{R} \cup \{-\infty\}$ has non-empty superdifferential at each point near which it is finite-valued.

• *The constraint set X_i* of agent i accounts for his (technological) restrictions. By assumption, that set is compact convex. Let bdX_i denote the *boundary* and $intX_i$ the *interior* of that set. Most real instances have, for some finite index set $J(i)$,

$$X_i = \{x_i \in \mathbb{X} : c_j(x_i) \geq 0 \text{ for every } j \in J(i)\}, \tag{1}$$

each $c_j : \mathbb{X} \to \mathbb{R}$ being concave, differentiable.

Anyway, an input profile $(x_i) \in \mathbb{X}^I$ is declared *feasible* if each $x_i \in X_i$ and $\sum_{i \in I} x_i = \sum_{i \in I} e_i$. We write $A := \left\{ (x_i) : \sum_{i \in I} x_i = \sum_{i \in I} e_i \right\}$ for the set of *allocations*.

• A *constraint qualification*, called the *Slater condition*, is convenient (but not essential): there exists a (strictly feasible) allocation $i \mapsto \hat{x}_i \in intX_i$. This condition ensures that each agent i gets (essential) *marginal profit*

$$M_i(x_i) := \partial \pi_i(x_i) - N_i(x_i), \tag{2}$$

with

$$N_i(x_i) := \{n \in \mathbb{X} : \langle n, \chi - x_i \rangle \leq 0 \text{ for all } \chi \in X_i\}$$

denoting the outward *normal cone* of X_i at x_i. Note that $N_i(x_i) = \{0\}$ at each $x_i \in intX_i$. Further, $N_i(x_i)$ is empty whenever $x_i \notin X_i$.

For instance (1), $N_i(x_i)$ is the convex cone generated by the antigradients $-c_j'(x_i)$, $j \in J(i)$, for which $c_j(x_i) = 0$. The Slater condition then means that $c_j(\hat{x}_i) > 0$ for all $j \in J(i)$.

• *Uncertainty* affects welfare. More precisely, agent i faces a random entity ψ_i the impact of which becomes clear only *after* x_i has been committed. Thus, i's *realized* payoff is a random amount, denoted $\Pi_i(x_i, \psi_i) \in \mathbb{R}$. Writing E for the expectation operator, we posit that $\Pi_i(x_i, \psi_i)$ be concave in x_i, integrable in ψ_i, and $\pi_i(x_i) := E\Pi_i(x_i, \psi_i)$.

Nothing precludes that some variates $\psi_i, i \in I$, coincide. It imports that their joint or marginal distributions may be unknown.

---

[3]For instance, when profit is a *reduced function* $\pi_i(x_i) := \max_y \tilde{\pi}(x_i, y)$, kinks and corners easily emerge. Lipschitz continuity is however, preserved.

Remarks on smoothness of data: One may think constraints more challenging, important and relevant than non-smooth objectives. The latter is, however, more basic. To see why, recall that π_i is assumed Lipschitz continuous with some modulus l_i; i.e., $|\pi_i(\chi) - \pi_i(\bar{\chi})| \le l_i \|\chi - \bar{\chi}\|$ for all $\chi, \bar{\chi} \in \mathbb{X}$. Let

$$d_i(\chi) := \min \{\|\chi - x_i\| : x_i \in X_i\}$$

denote the *distance* from $\chi \in \mathbb{X}$ to X_i and pick any constant $L_i > l_i$. Agent i may then just as well use the alternative objective $\pi_i - L_i d_i$ and completely ignore his constraint $x_i \in X_i$. Indeed, when $\chi \notin X_i$ has closest approximation (alias *projection*) $P_i\chi \in X_i$, any point $\bar{\chi}$ in the half-open segment $(\chi, P_i\chi]$ lies at distance $d_i(\bar{\chi}) = d_i(\chi) - \|\chi - \bar{\chi}\|$ from X_i. Therefore, $\pi_i(\bar{\chi}) \ge \pi_i(\chi) - l_i \|\chi - \bar{\chi}\| > \pi_i(\chi) - L_i \|\chi - \bar{\chi}\|$ whence

$$\pi_i(\bar{\chi}) - L_i d_i(\bar{\chi}) > \pi_i(\chi) - L_i \|\chi - \bar{\chi}\| - L_i \{d_i(\chi) - \|\chi - \bar{\chi}\|\} = \pi_i(\chi) - L_i d_i(\chi).$$

In other words: the term $L_i d_i$ serves as an *exact penalty function* for the constraint $x_i \in X_i$. That term totally alleviates the need to account separately for restrictions; see Clarke et al. (1998). There is, however, no escape from non-smooth data. In fact, d_i isn't differentiable at the boundary of X_i. So, "ignoring" constraints cause no loss of generality provided we use non-smooth functions to meter out "exact" penalties.

We record that the distance d_i is convex with *subdifferential* (Rockafellar, 1970)

$$\partial d_i(\chi) = \begin{cases} N_i(\chi) \cap B & \text{when } \chi \in X_i \\ (\chi - P_i\chi)/d_i(\chi) & \text{otherwise.} \end{cases} \tag{3}$$

Here B denotes the closed unit ball in $\mathbb{X}$, and P_i is the projection onto X_i.

3. EQUILIBRIUM

We emphasize that the environment is decomposable (i.e., free of externalities). Also, it qualifies as informationally decentralized. That is, knowledge about the triple (Π_i, X_i, e_i) remains quintessentially private, available merely to i, and maybe to a degree unrecorded. To complicate matters further, neither the distributions nor the realized values of ψ_i need be known by any party.

Suppose now that these imperfectly informed, self-interested agents $i \in I$ exchange resources. Their dealings are completely decentralized and voluntary. Posit that each good $g \in G$ be perfectly divisible and non-perishable. Our interest is first with price-supported equilibria:

Definition (Equilibrium). *A feasible allocation (x_i) is declared a **market equilibrium supported by price** p iff for all i*

$$p \in M_i(x_i) = \partial\pi_i(x_i) - N_i(x_i). \quad \square \tag{4}$$

In essence, (4) says the price should equal the marginal profit of each agent. To see this, suppose π_i differentiable at x_i. If $x_i \in int X_i$, (4) means the customary condition

$p = \pi_i'(x_i)$; that is, $p_g = \frac{\partial}{\partial x_{ig}} \pi_i(x_i)$ for each good g. Otherwise, when $x_i \in bdX_i$, a normal $n \in N_i(x_i)$ yields $n = \pi_i'(x_i) - p$ whence the *variational inequality*

$$\langle \pi_i'(x_i) - p, \chi - x_i \rangle \leq 0 \text{ for all } \chi \in X_i,$$

saying that no variation $\chi - x_i$ away from x_i is worthwhile under net price $\pi_i'(x_i) - p$.

More material on equilibrium is given in Appendix 1. Equilibrium is, of course, static in nature, displaying the (possibly boring) tranquillity of a steady state. Our interest is more with attainability of such distinguished outcomes than with their persistence. So we ask: *can some reasonable process eventually bring about equilibrium?* The next section displays an attractive process. For its implementation and statement *two* matters must first be taken care of:

1) *The normal cone $N_i(x_i)$ cannot practically be handled.* It is too large when $x_i \in bdX_i$, and it is empty when $x_i \notin X_i$. Replace the said cone therefore, with the truncated, globally defined, non-empty counterpart $L_i \partial d_i(x_i)$; see (3).

2) The other query is that the troublesome operator E in $\pi_i(x_i) := E\Pi_i(x_i, \psi_i)$ may make agent i *unable to compute his current "margin"* $\partial \pi_i(x_i)$. As said, to execute E is often "hard", be it mentally or numerically. And, of course, when the underlying distribution is unknown, no mean value is offhand computable. To alleviate these difficulties suppose agent i simply replaces $\partial \pi_i(x_i)$ with the "realized margin" $\partial \Pi_i(x_i, \psi_i)$, the partial differential being taken with respect to x_i.

After such replacements the preceding object $M_i(x_i)$ (2) gives way for the somewhat more "naive" but tractable version

$$\mathcal{M}_i(x_i, \psi_i) := \partial \Pi_i(x_i, \psi_i) - L_i \partial d_i(x_i). \tag{5}$$

This sort of realized margin becomes crucial next. Broadly, "gradient" $\gamma_i \in \mathcal{M}_i(x_i, \psi_i)$ will affect revision of x_i.

In our opinion, to pick at least one $\gamma_i \in \mathcal{M}_i(x_i, \psi_i)$ requires modest skill. Agent i must merely notice own marginal payoff $\partial \Pi_i(x_i, \psi_i)$ and the feasible choice $P_i x_i$ closest to x_i; see (3).

4. CONTINUOUS-TIME EXCHANGE

Under simple hypotheses about market efficiency - found say, in finance - prices and quantities should respond correctly and instantly to new information. Ample evidence and several studies tell though, that this need not happen. On the other hand, under stable conditions, experiments indicate that good approximation to market equilibrium obtains within reasonable time.

To reinforce the last point, and to prepare for subsequent modelling, this section verifies the stability of an auxiliary, purely theoretical, highly stylized process. A reader who takes stability on faith, may proceed directly to the next section.

In this part of the paper uncertainty is ignored, and time is continuous. At instant $t \geq 0$ the process to be modeled has just reached allocation $x(t) = [x_i(t)] \in A$. Suppose agent i finds it attractive there to move $x_i = x_i(t)$ in "direction"

$$D_i(x_i) := \partial [\pi_i - L_i d_i] (x_i). \tag{6}$$

The penalty term $L_i d_i$ reflects of course his concerns with maintaining or restoring $x_i \in X_i$. So, direction $D_i(x_i)$ mirrors incentive compatible, decentralized choice. Most likely though, an overall displacement $D(x) := [D_i(x_i)]$ would violate material balances. Therefore, $D(x)$ must possibly be bent, via a projection P_{T_A}, so as to become aligned with A. The upshot is that a *differential inclusion*

$$\dot{x} \in P_{T_A} D(x), \tag{7}$$

moves $x(t)$ in continuous time $t \geq 0$, (Aubin and Cellina, 1984). It starts at some $x(0) \in A$ and keeps $x(t) \in A$ thereafter. $\dot{x}$ is shorthand for $\frac{d}{dt} x(t)$. As said, the operator P_{T_A} projects $D(x)$ onto the tangent space $T_A = \left\{ (v_i) \in \mathbb{X}^I : \sum_{i \in I} v_i = 0 \right\}$ of A.

Operator P_{T_A}, by enforcing the aggregate resource constraint, doesn't quite square with our modelling philosophy. Indeed, who undertakes that operation? We shall see below that the said operator practically disappears. Until then, continuous-time dynamic (7) serves merely to identify where and how stability comes about.

For simplicity in stating that property let $X := \Pi_{i \in I} X_i$ be the product of individual domains. Denote by $n := |I|$ the number of players. Recall that π_i has Lipschitz modulus $l_i < L_i$. Let $\delta := \min_{i \in I}(L_i - l_i)$.

Theorem 1 (Asymptotic stability). • *From any initial allocation $x(0) \in A$, there emanates a unique, infinitely extendable solution $0 \leq t \mapsto x(t) \in A$ to system* (7).
• *When $x(0) \notin X$, that solution satisfies $x(t) \in X$ for all $t > \sqrt{n}d(x(0))/\delta$.*
• *If the state is feasible at some time, it stays feasible forever after. That is, the set X is absorbing.*
• *$x(t)$ converges to a price supported equilibrium.* □

5. DISCRETE-TIME REPEATED EXCHANGE

Theorem 1 inspires some faith in the attainability and stability of equilibrium. It motivates us to model repeated exchanges between random pairs of agents, operating within a stationary albeit stochastic setting.

For good reasons, the requirement $\sum_{i \in I} x_i = \sum_{i \in I} e_i$ must hold throughout; that is, $(x_i) \in A$ at each stage. Also, because $e_i \notin X_i$ isn't precluded, individual feasibility will not always be enforced. Ultimately however, agent i must settle on some $x_i \in X_i$.

The main part of the exchange mechanism functions as follows. Suppose agent i, while owning $x_i \in \mathbb{X}$, meets another agent j, who has x_j on his hand. They perceive then gradients $\gamma_i \in \mathcal{M}_i(x_i, \psi_i)$ and $\gamma_j \in \mathcal{M}_j(x_j, \psi_j)$ respectively; see (5). If commodity component $g \in G$ of $\gamma_i - \gamma_j$ is positive, agent i might convince his interlocutor j to transfer some amount of that good. In return j could receive another good for which the corresponding difference is negative. If this sort of bilateral barter amounts to scale $\gamma_i - \gamma_j$ by a factor $s > 0$, the two parties acquire updated holdings

$$x_i \leftarrow x_i + s(\gamma_i - \gamma_j) \quad \text{and} \quad x_j \leftarrow x_j + s(\gamma_j - \gamma_i), \tag{8}$$

respectively. Clearly, such updating preserves material balance; it never exits from the allocation space A.[4]

For supplementary interpretation of (8), suppose agent i actually regards γ_i as his *"private price"*. If so, he assigns value $v_i := \langle \gamma_i, \gamma_i - \gamma_j \rangle$ to the gradient difference. Similarly, agent j comes up with evaluation $v_j := \langle \gamma_j, \gamma_j - \gamma_i \rangle$. Since $v_i + v_j = \|\gamma_i - \gamma_j\|^2$, the value sum is positive whenever $\gamma_i \neq \gamma_j$. We shall say nothing on *value transfers* between traders. We require however, that a net *commodity transfer* $s(\gamma_i - \gamma_j)$ goes to agent i.

No trading platform is described. Nothing is specified as to whether or how agents bargain. It's left open who among $\{i, j\}$ proposes and who responds. In particular, the parties need neither identify nor report entities γ_i, γ_j. What imports is merely that both perceive the same difference $\gamma_i - \gamma_j$ and use it as guideline.

Update (8) happens with newly selected agents time and again. Instead of stating the procedure as an algorithm we prefer to depict it as an

Exchange process: It *starts* with *some stepsize* and allocation $(x_i) \in A$. It proceeds iteratively, in discrete time, with four events happening at *each stage*:

1) *Two producers i, j ($i \neq j$) then meet by chance. They are chosen independently,* in equiprobable manner.

2) *An independent pair (ψ_i, ψ_j) is sampled according to its fixed joint distribution.*

3) *"Gradients" $\gamma_i \in \mathcal{M}_i(x_i, \psi_i)$ and $\gamma_j \in \mathcal{M}_j(x_j, \psi_j)$ are selected.*

4) *Endowments are updated* by (8).

Note that multitudes of random factors impinge the exchanges from outside. Presence of noise often makes data smoothing necessary, this operation usually causing some delays. Here however, it will be made part of the process, as follows. The scale factor or *stepsize* $s = s_k > 0$, used at stage k, is updated subject to

$$\sum_{k=0}^{\infty} s_k = +\infty \ \text{ and } \ \sum_{k=0}^{\infty} s_k^2 < +\infty. \tag{9}$$

Broadly, the effect of (9) is to take moving averages. Thus agents will ultimately see mean values. Our main result, stated next, substantiates this:

Theorem 2 (Global asymptotic stability of price-supported market equilibrium). *Under (9) the stochastic allocation process generated by iterated bilateral exchanges, converges almost surely to a price-supported equilibrium.* $\square$

---

[4]Thus, as said before, the projection P_A onto the allocation space A becomes superfluous.

To interpret (9) suppose "time" $\tau_k := s_0 + \cdots + s_{k-1}$ has accumulated prior to stage k. The divergence condition in (9) simply says that $\tau_k \to +\infty$ whereas the convergence condition implies $\Delta\tau_k := \tau_{k+1} - \tau_k = s_k \to 0^+$. Broadly, numerical integration of (7) must proceed with stepsizes which dwindle but not too fast. While (9) constrains how the sequence (s_k) behaves *asymptotically*, it imposes no restrictions in the short or medium run. Accordingly, agents may, for long time, respond strongly to exchange opportunities.

Traders were matched randomly here. Numerous other protocols could govern dynamic matching (Gale, 2000). For example, agent i might encounter j periodically or only when $\|\gamma_i - \gamma_j\|$ is maximal. Qualitatively, all reasonable regimes give the same results.

6. ENTER UTILITY MAXIMIZERS

In this section each agent i derives utility $u_i(\pi_i(x_i))$ of payoff. To reflect non-satiation and decreasing margins we posit that $u_i : \mathbb{R} \to \mathbb{R}$ be concave and differentiable with $u_i' \geq$ some positive ε_i. Then the composite function $u_i \circ \pi_i$ is also concave with superdifferential $\partial u_i \circ \pi_i(x_i) = u_i'(\pi_i(x_i))\partial\pi_i(x_i)$.

It is natural now to declare an allocation (x_i) a market equilibrium supported by price p if for every i it holds that $u_i'(\pi_i(x_i))p \in \partial u_i \circ \pi_i(x_i) - N_i(x_i)$. But evidently, the latter inclusion is equivalent to (4).

Replace now (5) with

$$\mathcal{M}_i(x_i, \psi_i) := \partial\pi_i(x_i, \psi_i) - \frac{L_i}{u_i'(\pi_i(x_i))}\partial d_i(x_i)$$

and proceed exactly as before.

Theorem 3 (Convergence under non-transferable utility). *The stochastic process, generated by iterated exchanges, still converges almost surely to an equilibrium.* $\square$

Proposition 1 (Equilibrium and the representative agent). *Suppose there is a unique equilibrium $(\bar{x}_i)$ (as happens when all $u_i \circ \pi_i$ are strictly concave). Then that point solves the following planning problem of the representative agent*

$$\max \sum_{i \in I} \frac{1}{u_i'(\bar{x}_i)} u_i \circ \pi_i(x_i) \text{ subject to each } x_i \in X_i \text{ and } \sum_{i \in I} x_i = \sum_{i \in I} e_i. \ \square$$

7. AN ILLUSTRATION: THE KYOTO EMISSION MARKET

Exchange of rights to emit greenhouse gases is guided by the Kyoto Protocol (UN-FCCC, 1997). For simplicity, we consider only one good - the right to pollute - and there is no uncertainty besides matching.

Parameters are taken from Godal and Klaassen (2006) and originate from the computable general equilibrium model MERGE of Manne and Richels (1992).[5]

A global carbon tax ranging from 0 to 250 US dollars per ton carbon (USD/tC) was introduced in MERGE for the year 2010. The emission response was computed and used to estimate marginal payoff functions via OLS regression. Affine approximations were good for positive taxes. Since emissions are finite when no tax is imposed, marginal payoff functions take the form

$$\pi'_i(x_i) = \max\{0, a_i - b_i x_i\}$$

for all $i \in I$, where a_i, $b_i > 0$.

The endowment of permits for any Party to the Protocol is not specified as a physical quantity, only as a number to be multiplied with that party's 1990 emission. To fix endowments, the 1990 emission level given by MERGE were scaled by the factors listed in the Protocol (UNFCCC, 1997). The agents $i \in I$, and their triples (a_i, b_i, e_i) are presented in Table 1.

Table 1. Parameters. Marginal payoffs are measured in US dollars per ton carbon and endowments are in million tons carbon per year.

Agents	Payoff functions		Endowments
i	a_i	b_i	e_i
USA	1003	0.551	1251
OECD Europe	1883	1.813	860
Japan	1727	4.933	258
CANZ	693	2.216	251
EEFSU	1410	1.569	1314

CANZ is short for Canada, Australia and New Zealand, while EEFSU collects the countries in Eastern Europe and the Former Soviet Union. The parameters displayed imply a perfectly competitive equilibrium price with (without) the participation of the USA that equals 142.8 (0 respectively) USD/tC. All simulations presented below were terminated when the bid-ask spread $\max_i\left\{(\pi'_i(x_i))_{i\in I}\right\} - \min_i\{(\pi'_i(x_i))_{i\in I}\}$ was ≤ 5 USD/tC.

Simulations of exchange guided by the gradient difference: Trade was first modeled by adopting (8) with stepsizes $s_k = \frac{2}{1+k}$. A simulation featuring those countries that signed the Kyoto agreement in 1997 (including the USA) is given in

---

[5]MERGE (A **M**odel for **E**valuating the **R**egional and **G**lobal **E**ffects of Greenhouse Gas Reduction Policies) is an intertemporal computable general equilibrium model with a relatively detailed representation of the energy sector (a prime emitter of greenhouse gases). Only the energy-related CO_2 emissions were accounted for in the applied version, calibrated to the B2 scenario made for the Intergovernmental Panel on Climate Change (Nakicenovic et al., 2000). See also www.stanford.edu/group/MERGE/.
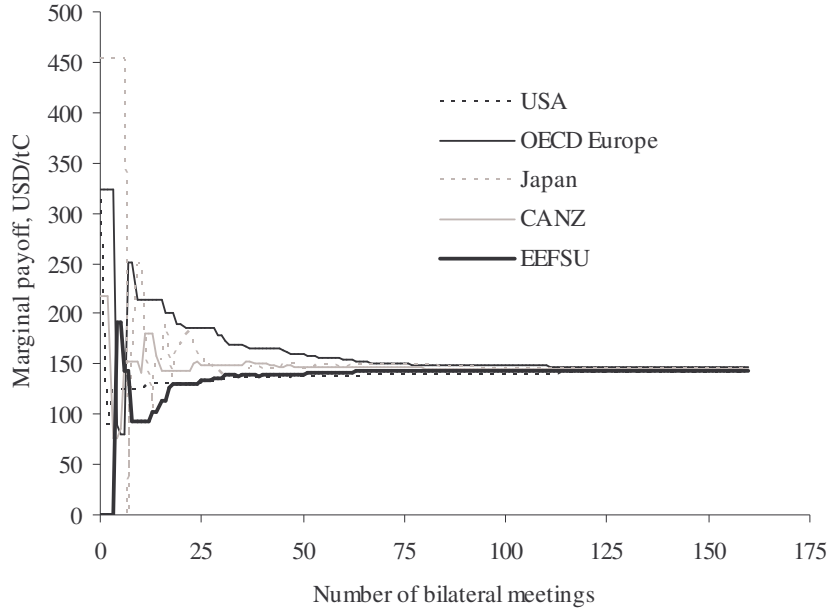
Figure 1.[6]

Figure 1: Bilateral permit exchange including the USA.

Figure 1 demonstrates that marginal payoffs converge to a competitive price. Since the USA hasn't yet ratified the Kyoto Protocol, the following picture better reflects the current state of affairs:

---

[6]The data behind Figures 1 to 4 are obviously discrete. The dots have been connected to improve presentation.
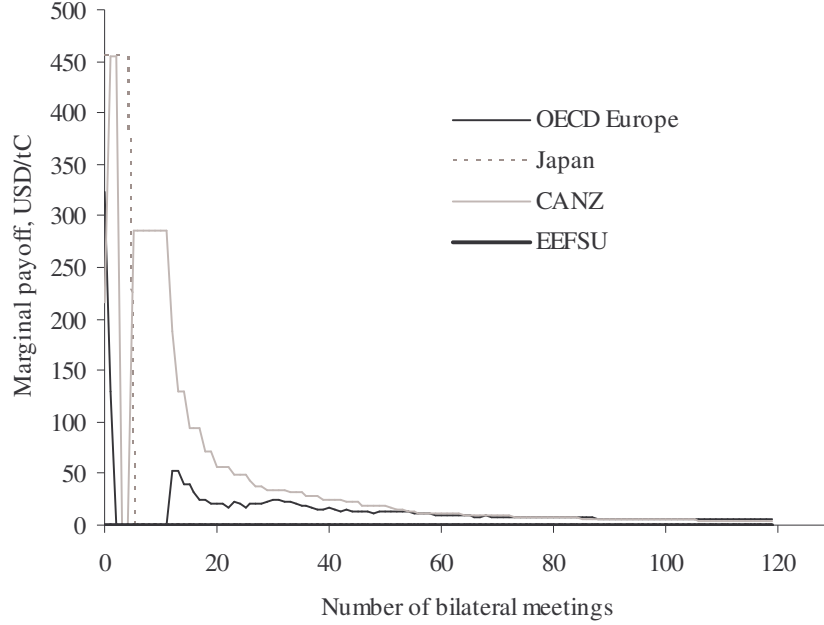
Figure 2: Bilateral permit exchange excluding the USA.

Without the USA, marginal payoffs are eventually nil because total supply exceeds aggregate demand. Specifically, $\sum_{i \in I \setminus \{USA\}} (e_i - a_i/b_i) = 2647 - 2600 > 0$, where a_i/b_i is commonly referred to as agent i's "business as usual emissions".

Alternative simulations: In place of (8) we now apply the normalized versions

$$x_i \leftarrow x_i + s \frac{\gamma_i - \gamma_j}{\|\gamma_i - \gamma_j\|} \quad \text{and} \quad x_j \leftarrow x_j + s \frac{\gamma_j - \gamma_i}{\|\gamma_j - \gamma_i\|}.$$

In the single good case, this alternative approach has the advantage of relieving traders from assessing gradient differences. It suffices that they identify the sign of that difference. The reason is that normalized forces, as studied in Ruszczyński (2006, Chap. 7), then amounts to using the sign function. With stepsizes $s_k = \frac{250}{1+k}$, $k = 0, 1, ..$ a simulation of trade with the USA is depicted in Figure 3,
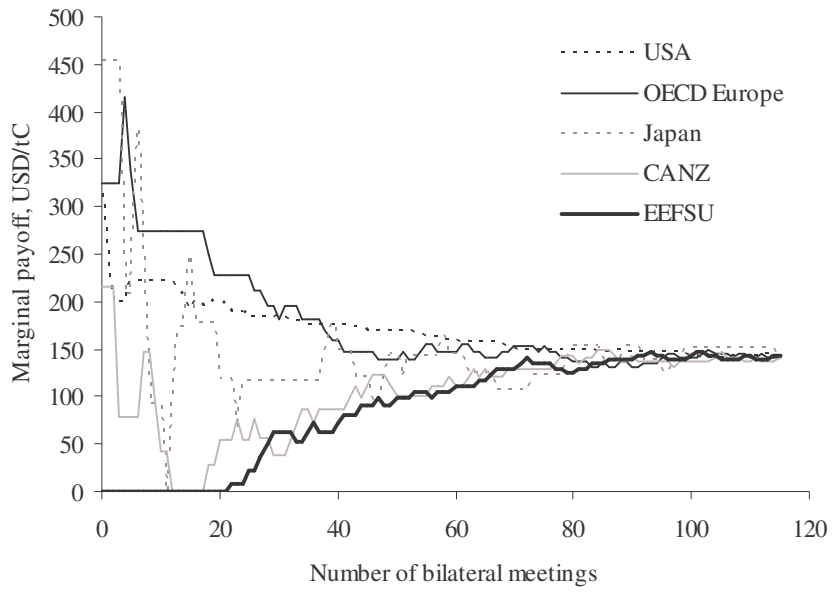
Figure 3: Bilateral permit exchange including the USA.

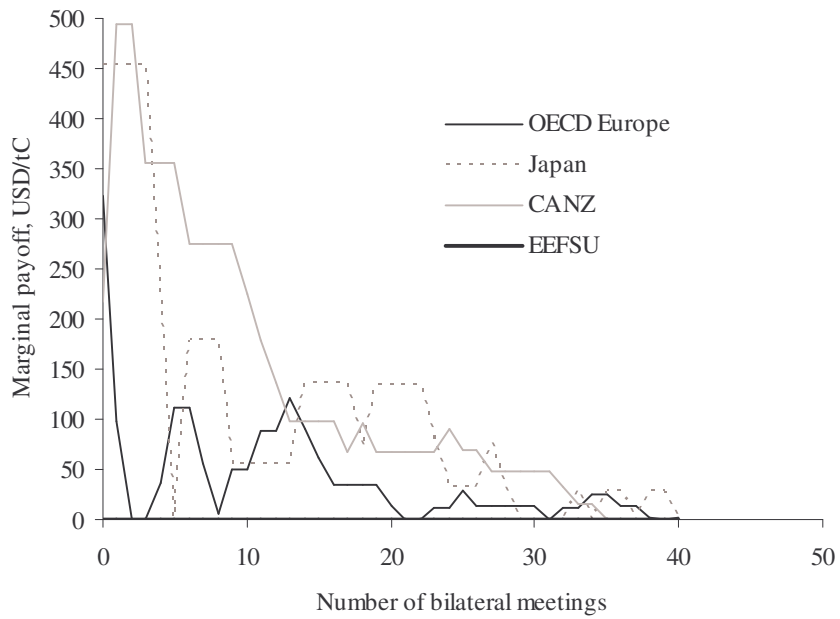while exchange excluding the USA is given in Figure 4.

Figure 4: Bilateral permit exchange excluding the USA.

Again, marginal payoffs converge to competitive prices.

8.   Concluding Remarks

Markets and prices occupy center stage in economic theory. Their functioning or formation is however, hard to model. It appears that some difficulties can be resolved by specifying how economic agents, while out of equilibrium, meet and trade. This paper opens up a simple vista on these things. It regards markets as recurrent encounters between potential traders. Each concerned agent holds a private resource bundle, and hesitates, at any stage, to reveal his marginal exchange rates. Also, he has limited and localized knowledge about the prospects for own improvement. This notwithstanding, any two meeting agents will detect whether a "small" transaction can benefit both. That is, they detect whether there are attractive prospects for a bilateral barter. Then, they exchange a moderate volume - most likely at implicit, idiosyncratic prices. In short: differences in substitution rates drive transactions. Granted transferable utility and decreasing margins, iterated exchange may eventually clear the market - and moreover, generate prices that support competitive equilibrium.

The paper differs fairly much from the received literature. No auctioneer or central agent directs or facilitates trade. Walrasian tâtonnement is therefore not an issue (Arrow and Hahn, 1971; Bala and Majumdar, 1992; Ermoliev et al., 2000a; Goree et al., 1998; Hurwitz et al., 1975a,b; Saari, 1985, 1995). Money might be present - and instrumental, say as a numeraire (Shapley and Shubik, 1977) - but its role is not explicit. Reliable prices will emerge therefore, only when all trade is completed and the market clears.

As usual, that outcome qualifies as core solution. We have however, not let coalitions play any role. Jevon's spirit dominates here in that pairwise interactions are main events (Ermoliev et al., 2000b). The flock of traders is fixed though; nobody arrives; nobody leaves. In that respect we are at variance with some models of sequential bargaining (Gale, 2000; Kunimoto and Serrano, 2004; McLennan and Sonnenschein, 1991; Rubinstein and Wolinsky, 1985, 1990). Most important, we dispense with the assumption that each party learns the type and current bundle of his interlocutor. Also different is that traders contend with "minor" transactions on each round. Transactions may occur that later are regretted. Such features of exchange have been observed in experiments; see Klaassen et al. (2005). Unlike Gintis (2006) no agent (or strategy) imitates, reproduces, or mutates. The paper that in parts comes closest to ours is Ermoliev et al. (2000b).

Appendix 1: On the Equilibrium

Collected here is material on *characterization/existence/uniqueness of equilibrium, core solutions, the inverse demand curve, the price slope, bid-ask spreads, normal cone representation, Lispschitz continuity, and interior point methods* - in that order. None of these issues interfere with the chief line of the paper.

Characterization, existence and uniqueness of equilibrium obtain easily from the hypotheses:

Proposition 2 (Characterization/existence/uniqueness of price-supported market equilibrium).
• *Given aggregate endowment $e_I := \sum_{i \in I} e_i$, an input profile (x_i) is a price-supported market equilibrium iff it solves the problem*

$$\pi_I(e_I) := max \left\{ \sum_{i \in I} \pi_i(x_i) : \sum_{i \in I} x_i = e_I \ and \ each \ \ x_i \in X_i \right\}. \qquad (10)$$

• *There exists at least one such equilibrium. It is unique if all π_i are strictly concave.*

Proof: For any set C let the extended indicator δ_C equal 0 on C and $+\infty$ elsewhere. Profile $x = (x_i)$ solves problem (10) iff

$$0 \in \partial \left\{ \sum_{i \in I} [\pi_i - \delta_{X_i}] - \delta_A \right\}(x),$$

where the total superdifferential ∂ is taken with respect to x. The constraint qualification ensures that

$$\partial \left\{ \sum_{i \in I} [\pi_i - \delta_i] - \delta_A \right\}(x) = [\partial \pi_i(x_i) - N_i(x_i)]_{i \in I} - N_A$$

where $N_A = \{(p, p, ...) : p \in \mathbb{X}\}$ is the normal cone to A at every $(x_i) \in A$. So, an allocation (x_i) is optimal iff there exists a price p such that (4) holds for each i. This takes care of the first bullet.

For the second bullet note that $(x_i) \mapsto \sum_{i \in I} \pi_i(x_i)$ is continuous on a non-empty compact domain. (In fact, the Slater profile $(\hat{x}_i)$ is feasible.) Existence of a maximum is then guaranteed. The assertion about uniqueness derives from the strict concavity of the objectives. $\square$

Equilibrium as a core solution to a cooperative game:

Proposition 3 (Equilibrium as a core solution, Evstigneev and Flåm, 2001). *Define a transferable-utility game with player set I and characteristic function*

$$I \supseteq S \mapsto \pi_S(e_S) := \sup \left\{ \sum_{i \in S} \pi_i(x_i) : \ x_i \in X_i \ \ and \ \sum_{i \in S} x_i = \sum_{i \in S} e_i =: e_S \right\}.$$

Then, each equilibrium price p defines a core solution $c = (c_i) \in \mathbb{R}^I$ by

$$c_i := \sup \left\{ \pi_i(\chi) - \langle p, \chi \rangle : \chi \in X_i \right\} + \langle p, e_i \rangle$$

in that

$$\sum_{i \in S} c_i \geq \pi_S(e_S) \ \text{for all } S \subseteq I \text{ with equality for } S = I. \ \square$$

In fact, $\sum_{i \in S} c_i \geq \pi_S(e_S)$ for *any* price vector p. This inequality mirrors that access to a competitive market can harm no coalition $S \subseteq I$. To preclude $\sum_{i \in I} c_i > \pi_I(e_I)$ an equilibrium price must be used, meaning a p for which

$$\pi_I(e_I) \geq \sum_{i \in I} \sup \left\{ \pi_i(x_i) + \langle p, e_i - x_i \rangle : x_i \in X_i \right\}.$$

The "price curve" $p = P(\sum_{i \in I} e_i)$ satisfies the "*law of demand*":

Proposition 4 (Properties of the derived price curve).
- *At least one price supports a market equilibrium.*
- *A price applies for such support iff it belongs to $\partial \pi_I(e_I)$.*
- *In particular, the price varies only with the aggregate supply $e_I = \sum_{i \in I} e_i$.*
- *The price curve $P(\cdot) := \partial \pi_I(\cdot)$ "slopes downwards":*

$$\langle p - \bar{p}, e_I - \bar{e}_I \rangle \leq 0 \ \text{for all } p \in P(e_I) \ \ and \ \ \bar{p} \in P(\bar{e}_I). \tag{11}$$

- *If some function $p \mapsto \sup \left\{ \pi_i(x_i) - \langle p, x_i \rangle : x_i \in X_i \right\}$ is strictly convex, then $P(e_I)$ is unique and continuous.*

Proof. The function π_I defined in (10) is concave. Moreover, by the constraint qualification, π_I is finite-valued near the given aggregate e_I.[7] For that reason, there exists at least one $p \in \partial \pi_I(e_I)$. By Proposition 2 some allocation (x_i) solves problem (10). Consequently, $p \in \partial [\pi_i - \delta_{X_i}] (x_i) = \partial \pi_i(x_i) - N_i(x_i)$ for all i; see Laurent (1972). The concavity of π_I implies that $\partial \pi_I$ must be monotone decreasing as expressed in (11).

Each function $f : \mathbb{X} \to \mathbb{R} \cup \{-\infty\}$ has a *conjugate*

$$f^{(*)}(x^*) := \sup \left\{ f(x) - \langle x^*, x \rangle : x \in \mathbb{X} \right\}.$$

---

[7]It suffices for having $\pi_I(\cdot)$ finite near e_I that some feasible allocation $(\hat{x}_i)$ has $\hat{x}_i \in int X_i$ for at least one i.

Because $\pi_I^*(p) = \sum_{i \in I} [\pi_i - \delta_{X_i}]^{(*)}(p)$, with at least one term strictly convex, the final assertion follows from Theorem 4.1.1, page 79 in Hiriart-Urruty and Lemaréchal (1993). $\square$

The slope of inverse demand sometimes obtains as follows. Suppose $\pi_I(e_I)$ is uniquely attained by a feasible allocation (x_i). (This happens for instance if all π_i are strictly concave.) Then, under suitable regularity conditions, the price curve $P(\cdot)$ is differentiable at e_I with "slope"

$$P'(e_I) = \left\{ \sum_{i \in I} \pi_i''(x_i)^{-1} \right\}^{-1}, \tag{12}$$

see Crouzeix (1977). Like each $\pi_i''(x_i)$ the "slope" $P'(e_I)$ is a negative definite $G \times G$ Jacobian matrix. Interpret $\pi_i''(x_i)$ as a sort of "*resistance*." Accordingly, $\pi_i''(x_i)^{-1}$ reflects corresponding "*conductance*." Formula (12) points to electrical engineering in saying that the conductance of a parallel circuit equals the sum of conductances (Dorf and Svoboda, 1999). In general, we cannot expect that $P(\cdot)$ be differentiable. It appears though, that generalized, second-order derivatives are applicable (Rockafellar and Wets, 1998).

Bid-ask spreads may occur as illustrated next. At $x_i \in int X_i$ agent i would *bid*

$$\beta_i \leq \pi_i'(x_i; d) := \lim_{s \to 0^+} \frac{\pi_i(x_i + sd) - \pi_i(x_i)}{s}$$

for additional resources in direction d, $\|d\| = 1$. Similarly, he would *ask* $\alpha_i \geq -\pi_i'(x_i; -d)$ for giving away resources along that direction. Then, by concavity,

$$\beta_i \leq \pi_i'(x_i; d) \leq -\pi_i'(x_i; -d) \leq \alpha_i,$$

the middle inequality being strict when $\pi_i(\cdot)$ has a kink at x_i along direction d.

The normal cone $N_i(x_i)$ at $x_i \in X_i$, when X_i is given by (1), consists of all vectors $-\sum_{j \in J(i)} \lambda_j c_j'(x_i)$ with each $\lambda_j \geq 0$, and $\lambda_j c_j(x_i) = 0$.

Lipschitz continuity of π_i, if absent, could be ensured as follows: Take any positive number r_i and replace π_i by the regularized counterpart

$$\hat{\pi}_i(x_i) := \sup \left\{ \pi_i(\chi) - r_i \|\chi - x_i\|^2 : \chi \in X_i \right\},$$

embodying adjustment costs $r_i \|\chi - x_i\|^2$. That substitute function $\hat{\pi}_i$ is bounded above, concave, Fréchet differentiable and Lipschitz continuous on the entire space; see Theorem 1.5.1 in Clarke et al. (1998).

Interior point methods offer alternative procedures to any agent $i \in I$ whose domain X_i has non-empty interior. That agent could add a *barrier term* $b_i(x_i)$ to $\pi_i(x_i)$ to maintain own feasibility. Specifically, the *barrier function* b_i, defined and concave on $int X_i$, should tend towards $-\infty$ whenever $x_i \in int X_i$ approaches $bd X_i$. For example, if $X_i = [0, 1]$, the function $b_i(x_i) = \varepsilon_i \{\log x_i + \log(1 - x_i)\}$ would do for any $\varepsilon_i > 0$. This approach has the advantage of dispensing with the distance function, but the drawback of merely producing approximate equilibria.

Closer scrutiny shows that our method applies even if $int X_i$ is empty. It requires merely that $\mathcal{M}_i$ (5) be replaced by

$$\mathbb{M}_i(x_i, \psi_i) := \partial \left[\Pi_i(\cdot, \psi_i) - L_i d_i \right](x_i).$$

Appendix 2: Proofs

Proof of Theorem 1. • (*For existence and uniqueness of solutions*) Note that $D(x)$ is *monotone decreasing*, meaning

$$\langle D(x) - D(\bar{x}), x - \bar{x}\rangle := \sum_{i \in I} \langle D_i(x_i) - D_i(\bar{x}_i), x_i - \bar{x}_i\rangle \leq 0$$

for any $x = (x_i), \bar{x} = (\bar{x}_i) \in \mathbb{X}^I$. In fact, monotonicity derives from $\pi_i - L_i d_i$ being concave. Consequently, its "derivative" $D_i = \partial\left[\pi_i - L_i d_i\right]$ decreases:

$$\langle D_i(x_i) - D_i(\bar{x}_i), x_i - \bar{x}_i\rangle \leq 0 \text{ for any } x_i, \bar{x}_i \in \mathbb{X}.$$

By an infinitely extendable *solution* is understood an absolutely continuous function $x(\cdot) : \mathbb{R}_+ \to A$ that satisfies (7) for almost every t. It is known that solutions coincide with those of $\dot{x} \in D(x) - N_A(x)$ where $N_A(x)$ denotes the normal cone to A at $x \in A$. The latter system remains monotone - in fact, *maximal* as such - hence has a unique, infinitely extendable solution; see Aubin and Cellina (1984). This takes care of the first bullet.

• (*For feasibility in finite time*) Endow $\mathbb{X}^I$ with the natural inner product $\langle x, \bar{x}\rangle := \sum_{i \in I}\langle x_i, \bar{x}_i\rangle$ and associated norm $\|\cdot\|$. The distance $d(x) := \|x - \bar{x}\|$ from a point $x \notin X$ to its closest approximation $\bar{x} \in X$ has derivative $d'(x) = (x - \bar{x})/d(x)$. Let $I(x) := \{i \in I : x_i \in X_i\}$. While $x(t) \notin X$, consider the function

$$\lambda(t) := d^2(x(t))/2 = \|x(t) - \bar{x}(t)\|^2/2. \tag{13}$$

During that phase, omit repeated mention of time to get

$$
\begin{aligned}
\dot{\lambda} &= d(x)\langle d'(x), \dot{x}\rangle = \langle x - \bar{x}, \dot{x}\rangle \in \langle x - \bar{x}, D(x) - N_A(x)\rangle \\
&\leq \langle x - \bar{x}, D(x)\rangle = \sum_{i \notin I(x)}\left\langle x_i - \bar{x}_i, \partial\pi_i(x_i) - L_i\frac{x_i - \bar{x}_i}{d_i(x_i)}\right\rangle \\
&= \sum_{i \notin I(x)}\left\{\langle x_i - \bar{x}_i, \partial\pi_i(x_i)\rangle - L_i d_i(x_i)\right\} \quad (*) \\
&\leq \sum_{i \in I}(l_i - L_i)d_i(x_i) \leq -\delta\sum_{i \in I}d_i(x_i) \leq -\delta\max_{i \in I}d_i(x_i) \leq -\frac{\delta}{\sqrt{n}}d(x).
\end{aligned}
$$

In (*) we used the Cauchy-Schwarz inequality and the fact that each supergradient in $\partial\pi_i(x_i)$ is bounded in norm by the Lipschitz constant l_i of π_i. The very last inequality follows from

$$\max_{i \in I}d_i(x_i) = \frac{1}{\sqrt{n}}(n\max_{i \in I}d_i^2(x_i))^{1/2} \geq \frac{1}{\sqrt{n}}(\sum_{i \in I}d_i^2(x_i))^{1/2} = \frac{1}{\sqrt{n}}d(x).$$

Because $\dot{\lambda} = d(x)\dot{d}(x)$ the proven inequality $\dot{\lambda} \leq -\delta d(x)/\sqrt{n}$ implies $\dot{d}(x(t)) \leq -\delta/\sqrt{n}$ while $x(t) \notin X$. During that phase

$$d(x(t)) = d(x(0)) + \int_0^t \dot{d} \leq d(x(0)) - t\delta/\sqrt{n} \tag{14}$$

whence $t \leq \sqrt{n}d(x(0))/\delta$. This proves the second bullet.

• (*Absorption*) Suppose x leaves X at some time. Then reset the clock to 0. From (14) follows that $d(x(t))$ immediately thereafter goes negative - an absurdity.

• (*For convergence to equilibrium*) Consider next the "ultimate" phase during which feasibility prevails all the time. In that regime let $\bar{x}(t)$ denote the *equilibrium* closest to $x(t)$ and consider again a function λ of the form (13). Arguing as above we get

$$\dot{\lambda} \leq \sum_{i \in I} \langle x_i - \bar{x}_i, \partial\left[\pi_i - L_i d_i\right](x_i)\rangle \leq \sum_{i \in I} \{x_i - \bar{x}_i, \partial\left[\pi_i - L_i d_i\right](\bar{x}_i)\}.$$

Since $\{x_i - \bar{x}_i, \partial\left[\pi_i - L_i d_i\right](\bar{x}_i)\} \leq 0$ for each i, we get $\dot{\lambda} \leq 0$. Strict inequality holds while $x \neq \bar{x}$. Indeed, then $x_i \neq \bar{x}_i$ for some i, and

$$\langle x_i - \bar{x}_i, \partial\left[\pi_i - L_i d_i\right](x_i)\rangle < \pi_i(x_i) - \pi_i(\bar{x}_i) < 0.$$

This completes the proof. $\square$

Proof of Theorem 2. Let $n \geq 2$ denote the number of agents. Under equiprobable choice of agent pairs, the chance of selecting any specific pair (i,j) is $\mu := 1/\binom{n}{2} = 2/\{n(n-1)\}$.

Write $\psi = (\psi_i)$ to indicate the exogenous state of the world. Consider the extended event space Ω composed of elementary outcomes $\omega = (i,j,\psi)$. Each such triple features two distinct agents i,j and the exogenous state ψ. Endow Ω with the product sigma-algebra and the corresponding probability measure. That is, for any measurable outside event Ψ, composed of elementary exogenous outcomes ψ, and for any agent pair (i,j) of distinct players, let

$$\Pr\left[(i,j) \times \Psi\right] = \mu \Pr\left[\Psi\right]$$

be the associated probability. Here $\Pr\left[\Psi\right]$ is the time-invariant probability already assigned to the set Ψ.

Consider any stage $k \geq 0$. For simpler notation suppress mention of k. What enters that stage is stepsize s, prevailing profile $x = (x_i)$, random event $\omega = (i,j,\psi)$, and gradients $\gamma_i \in \mathcal{M}_i(x_i,\psi_i)$, $\gamma_j \in \mathcal{M}_j(x_j,\psi_j)$. Define a new profile x^{+1} by

$$\left.\begin{array}{rcl} x_i^{+1} & := & x_i + s(\gamma_i - \gamma_j) \\ x_j^{+1} & := & x_j + s(\gamma_j - \gamma_i) \\ x_{\mathbf{i}}^{+1} & := & x_{\mathbf{i}} \text{ when } \mathbf{i} \neq i,j \end{array}\right\} \tag{15}$$

Expectation with respect to ψ_i in (5) yields

$$EM_i(x_i, \psi_i) = \partial \left[\pi_i - L_i d_i\right](x_i) = D_i(x_i).$$

Taking expectation with respect to (i, j) as well in (15) gives

$$Ex_i^{+1} \in x_i + \mu s \sum_{j \in I} \{D_i(x_i) - D_j(x_j)\} \text{ for each } i \in I. \tag{16}$$

For any non-empty closed convex set C, let P_C be the orthogonal projection operator onto that set. Given any allocation $x \in A$ and vector $v \in \mathbb{X}^I$, we have $P_A[x + snv] = x + s\left[\sum_{j \in I}\{v_i - v_j\}\right]_{i \in I}$. Thus, in terms of the *tangent cone* $T_A = \left\{(x_i) \in \mathbb{X}^I : \sum_{i \in I} x_i = 0\right\}$ of the allocation space A we get

$$\lim_{s \to 0^+} \frac{P_A[x + snv] - x}{s} = \frac{P_A[x + snv] - x}{s} = \left[\sum_{j \in I}\{v_i - v_j\}\right]_{i \in I} = P_{T_A}[nv].$$

As before $D(x) := [D_i(x_i)]_{i \in I}$. Substituting $D(x)$ for v in the last string gives

$$\frac{1}{n}\left[\sum_{j \in I}\{D_i(x_i) - D_j(x_j)\}\right]_{i \in I} = P_{T_A}[D(x)].$$

Thus, at stage k (16) assumes the form

$$Ex^{k+1} \in x^k + s_k n \mu P_{T_A}\left[D(x^k)\right],$$

showing that exchange amounts - in expectation - to an explicit Euler step of size $s_k n \mu$ of the differential inclusion

$$\dot{x} \in n \mu P_{T_A}[D(x)]. \tag{17}$$

Clearly (17) is just a scaled version of (7), inheriting the same qualitative properties.

Having now related iterated exchange via (16) to the asymptotically stable system (17), convergence of the exchange process follows from received results in stochastic approximation theory; see Benaim (1996) and Benveniste et al. (1990). This completes the proof. $\square$

Proof of Theorem 3. The demonstration of Theorem 2 also applies here. We provide however, an alternative demonstration (applicable for Theorem 2 as well). Let $\bar{x}^k$ be the optimal solution to (10) closest to x^k and define $a_k := \left\|x^k - \bar{x}^k\right\|^2$. When the event $\omega^k = (i^k, j^k, \psi^k)$ is sampled at stage k, use (15) to define

$$x_i^{k+1} := x_i^k + s_k v_i^k \text{ for all } i.$$

Here $v_i^k = v_i^k(x_i^k, \omega^k) \in \mathcal{M}_i(x_i, \psi_i) - \mathcal{M}_j(x_j, \psi_j)$ if $\{i, j\} = \{i^k, j^k\}$, 0 otherwise, and $\sum_{i \in I} v_i^k = 0$. Clearly,

$$a_{k+1} \leq \left\| x^{k+1} - \bar{x}^k \right\|^2 = \left\| x^k - \bar{x}^k + s_k v^k \right\|^2 = a_k + 2s_k \left\langle v^k, x^k - \bar{x}^k \right\rangle + s_k^2 \left\| v^k \right\|^2. \quad (18)$$

In terms of the sigma-field $\mathcal{F}_k$ generated by $\omega^0, ..., \omega^{k-1}$ take conditional expectations to define

$$b_k := E\left[\left\langle v^k, \bar{x}^k - x^k \right\rangle | \mathcal{F}_k \right] = \left\langle E\left[v^k | \mathcal{F}_k \right], \bar{x}^k - x^k \right\rangle \geq 0 \ \text{ and } \ c_k := E\left[\left\| v^k \right\|^2 | \mathcal{F}_k \right].$$

Since $c_k \leq C(1 + a_k)$ for some constant C, taking the same conditional expectation $E\left[\cdot | \mathcal{F}_k \right]$ through (18) gives

$$E\left[a_{k+1} | \mathcal{F}_k \right] \leq (1 + Cs_k^2)a_k - 2s_k b_k + s_k^2 C.$$

The assumption $\sum s_k^2 < +\infty$ implies that $a_k \to a$ and $\sum_k s_k b_k < +\infty$ almost surely; see Chap. 5 in Benveniste et al. (1990). We claim that $a = 0$ a.s. Indeed, whenever $a \neq 0$ along some scenario, b_k is - along that same scenario - eventually bounded away from zero. Then $\sum s_k = +\infty$ entails the contradiction $\sum_k s_k b_k = +\infty$. $\square$

References

Arrow, K.J., Hahn, F.H., 1971. General Competitive Analysis. Holden-Day, San Fransisco.

Aubin, J.P., Cellina, A., 1984. Differential Inclusions. Springer, Berlin.

Bala, B., Majumdar, M., 1992. Chaotic tâtonnement. Economic Theory 2, 437-445.

Benaim, M., 1996. A dynamical system approach to stochastic approximations. SIAM Journal on Control and Optimization 43, 437-472.

Benveniste, A., Métivier, M., Priouret, P., 1990. Adaptive Algorithms and Stochastic Approximations. Springer-Verlag, Berlin.

Clarke, F.H., Ledyaev, Yu.S., Stern, R.J., Wolenski, P.R., 1998. Nonsmooth Analysis and Control Theory. Springer, Berlin.

Crouzeix, J.-P., 1977. A relationship between the second derivative of a convex function and of its conjugate. Mathematical Programming 13, 364-365.

Dorf, R.C., Svoboda, J.A., 1999. Introduction to Electric Circuits. J. Wiley & Sons, New York.

Ermoliev, Y., Wets, R.J.-B., 1988. Numerical Techniques in Stochastic Programming. Springer, Berlin.

Ermoliev, Y., Keyzer, M.A., Norkin, V., 2000a. Global convergence of the stochastic tâtonnement process. Journal of Mathematical Economics 34, 173-190.

Ermoliev, Y., Michalevich, M., Nentjes, A., 2000b. Markets for tradeable emission and ambient permits: a dynamic approach. Environmental and Resource Economics 15, 39-56.

Evstigneev, I.V., Flåm, S.D., 2001. Sharing nonconvex costs. Journal of Global Optimization 20, 257-271.

Gale, D., 2000. Strategic Foundations of General Equilibrium. Cambridge University Press, Cambridge.

Gintis, H., 2006. The emergence of a price system from decentralized bilateral exchange. Contributions to Theoretical Economics 6(1), Art. 13.

Godal, O., Klaassen, G., 2006. Carbon trading across sources and periods constrained by the Marrakesh Accords. Journal of Environmental Economics and Management 51, 308-322.

Gode, D.K., Sunder, S., 1993. Allocative efficiency of markets with zero intelligence traders: markets as a partial substitute for individual rationality. Journal of Political Economy 101, 119-37.

Goeree, J.K., Hommes, C., Weddepohl, C., 1998. Stability and complex dynamics in a discrete tâtonnement model. Journal of Economic Behavior and Organization 33, 395-410.

Grossman, S.J., Stiglitz, J.E., 1976. Information and competitive price systems. American Economic Review 66, 246-253.

Hayek, F.A., 1945. The use of knowledge in society. American Economic Review 35, 519-530.

Hiriart-Urruty J.-B., Lemaréchal, C., 1993. Convex Analysis and Minimization Algorithms II. Springer-Verlag, Berlin.

Hurwicz, L., Radner, R., Reiter, S., 1975a. A stochastic decentralized resource allocation process: Part I. Econometrica 43, 187-221.

Hurwicz, L., Radner, R., Reiter, S., 1975b. A stochastic decentralized resource allocation process: Part II. Econometrica 43, 363-393.

Klaassen, G., Nentjes, A., Smith, M., 2005. Testing the theory of emissions trading: Experimental evidence on alternative mechanisms for global carbon trading. Ecological Economics 53, 47-58.

Kunimoto, T., Serrano, R., 2004. Bargaining and competition revisited. Journal of Economic Theory 115, 78-88.

Laurent, P.-J., 1972. Approximation et optimisation. Hermann, Paris.

Manne, A.S., Richels, R.G., 1992. Buying Greenhouse Insurance: the economic costs of carbon dioxide emission limits. MIT Press, Cambridge, MA.

McLennan A., Sonnenschein, H., 1991. Sequential bargaining as a noncooperative

foundation of competitive equilibrium. Econometrica 59, 1395-1421.

Nakicenovic, N., Alcamo, J., Davis, G., de Vries, B., Fenhann, J., et al., 2000. Emissions Scenarios, A Special Report of Working Group III of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge.

Rockafellar, R.T., 1970. Convex Analysis. Princeton University Press, New Jersey.

Rockafellar, R.T., Wets, R.J.-B., 1998. Variational Analysis. Springer-Verlag, Berlin.

Rubinstein, A., Wolinsky, A., 1985. Equilibrium in a market with sequential bargaining. Econometrica 53, 1133-1150.

Rubinstein, A., Wolinsky, A., 1990. Decentralized trading, strategic behavior and the Walrasian outcome. Review of Economic Studies 57, 63-78.

Ruszczyński, A., 2006. Nonlinear Optimization. Princeton University Press, Princeton.

Saari, D.G., 1985. Iterative price mechanisms. Econometrica 53, 1117-1131.

Saari, D.G., 1995. Mathematical complexity of simple economics. Notes of the American Mathematical Society 42, 222-230.

Shapley, L., Shubik, M., 1977. Trade using one commodity as a means of payment. Journal of Political Economy 85, 937-968.

Smith, V.L., 1962. An experimental study of competitive market behavior. Journal of Political Economy 70, 11-37.

Smith, V.L., 1976. Experimental economics: induced value theory. American Economic Review 66, 274-279.

Smith, V.L., 1982. Microeconomic systems as an experimental science. American Economic Review 72, 923-955.

Tesfatsion, L., Judd, K.L., 2006. Handbook of Computational Economics II: Agent-Based Computational Economics. North-Holland.

UNFCCC, 1997. Kyoto Protocol to the United Nations framework convention on climate change. Report of the Conference of the Parties, Third Session, Kyoto (available at http://unfccc.int/).

Department of Economics
University of Bergen
Fosswinckels gate 6
N-5007 Bergen, Norway
Phone: +47 55 58 92 00
Telefax: +47 55 58 92 10
http://www.svf.uib.no/econ