

THE COMPLEXITY OF EXCHANGE*

Robert Axtell

The computational complexity of two classes of market mechanisms is compared. First the Walrasian interpretation in which prices are centrally computed by an auctioneer. Recent results on the *computational complexity* are reviewed. The *non-polynomial complexity* of these algorithms makes Walrasian general equilibrium an *implausible* conception. Second, a *decentralised* picture of market processes is described, involving concurrent exchange within transient coalitions of agents. These processes feature *price dispersion*, yield allocations that are *not in the core*, modify the *distribution of wealth*, are always *stable*, but *path-dependent*. Replacing the Walrasian framing of markets requires substantial revision of conventional wisdom concerning markets.

1. Markets and the Emergence of Prices

Consider the following strategic environment. There is a heterogeneous population of autonomous entities, each of whom has internal states that describe its self-interest as well as certain external states. Each entity is engaged in purposive activity to further its interests, including altering its external state in exchange for alterations in the external states of other agents. Each individual receives information from other individuals directly, and has access to some global state information as well, although no agent has complete information on the global state. Calling these entities *agents*, we imagine that each one engages in more or less *strategic behaviour*. That is, each agent has some internal model for how the individuals in the population will behave and uses this model in order to decide how best to act subsequently. Finally, there exist performance measures, both subjective and objective, for the individuals as well as for the overall system of agents.

At a very high level of abstraction, this picture of interacting agents can describe a great variety of human activity, economic activity in particular. It might be a story about consumption behaviour insofar as the agents are humans who exchange money for goods, making decisions at least partly on the basis of information – on product quality, say – received from others. It could also be a framework for studying the operation of firms, in which individual actions are sufficiently coordinated that economic goods result from the interactions. Here, agents must communicate the nature of their productive actions to their peers and adapt their behaviour as external data arrives and their internal models are updated. This abstract depiction of interacting agents can also be a model of markets, in which

* This work has benefited from conversations with Ann Bell, Joshua Epstein, Doyne Farmer, Alan Kirman, Axel Leijonhufvud, Kai Nagel, Kishlaya Prasad, John Rust, Nick Vriend, Michael Wellman, and Peyton Young, from seminar participants at Brookings, CalTech, the European University Institute, Johns Hopkins, UCLA, Virginia Tech and Yale, and from workshop and conference participants at Autonomous Agents, the Society for Computational Economics and the Santa Fe Institute. Financial support from the John D. and Catherine T. MacArthur Foundation and the National Science Foundation (grant IRI-9820872) is gratefully acknowledged.

the agents exchange items of value about which each individual has its own, typically private, assessments. Such private valuations may *not* depend significantly on how other agents value the good, as when the item provides a service from which the agent benefits, e.g., the transportation service a car provides. Alternatively, the items being traded may have value to a particular agent that depends in an essential way on how other agents value it, e.g., the resale value of a car. This is the case of financial markets, in which the items being exchanged provide little utility intrinsically but rather have value insofar as they can be exchanged for consumption goods at later times.

Beyond purely economic activity, this abstract conception of interacting, self-interested entities is also a credible portrait of other complex adaptive systems, both natural and artificial. In an ant colony each individual ant performs a function, using local, socially-transmitted information in order to do its task. In the human immune system heterogeneous cell types interact to synthesise antibodies in order to neutralise invaders. In engineered systems, like computer networks, it is increasingly common for individual nodes to have well-defined objective functions (e.g., keep busy) and a behavioural repertoire that attempts to further that objective. Indeed, consider the general problem of distributed computing in which a single task is divided into pieces so that it can be worked on by several computers at once. Each individual computer in the network works on its piece of the larger task while communicating intermediate results to other computers. In such circumstances each node may locally adapt its behaviour as its instantaneous duties change, attempting to achieve a balanced load across the population of processors for instance (Bertsekas and Tsitsiklis, 1989). Alternatively, the load might be balanced from the top down, perhaps by a dedicated processor that is otherwise off-line, not part of the main distributed computation. Performance measures for such systems increase in the speed of obtaining a solution to the problem, a metric that is typically an increasing function of the quality of load balancing achieved.

This paper is primarily concerned with drawing out the connections between economic exchange, on the one hand, and distributed computation on the other, linked through this abstract picture of interacting, purposive agents. We shall argue that actions by self-interested agents in economic markets have much in common with the decentralised interactions of processors in distributed computation environments (Brewer, 1999; Cheng and Wellman, 1998). Using the asynchronous model of distributed computing, we shall study the performance of market systems as a function of their scale, i.e., the number of agents in the marketplace and the number of commodities being traded. The computational complexity of such systems is analysed and compared to that of conventional market models having centralised price determination, i.e., the Walras-Arrow-Debreu model.

A more practical motivation of this paper arises from a class of market models known as 'agent-based artificial markets'. In agent-based computational models a population of software objects is instantiated and each agent is given certain internal states (e.g., preferences, endowments) and rules of behaviour (e.g., seek utility improvements) (Epstein and Axtell, 1996; Delli Gatti *et al.*, 2000). The agents

are then permitted to interact directly with one another and a macrostructure emerges from these interactions. Patterns in this macrostructure may then be compared with empirical data, agent internal states and rules revised, and the process repeated until an empirically plausible model obtains. Models of this type are capable of reproducing both stylised features of financial markets (Arthur *et al.* (1997), Chen and Yeh (1997) and LeBaron (2001*a-c*), as well as many quantitative facts (Lux 1998; Levy *et al.*, 2000; Darley *et al.*, 2001; Zawadowski *et al.*, 2002).

An important aspect of these agent market models is the price formation process.¹ Through the interactions of the agents prices *emerge* in such models, either mediated by middlemen or market makers or not. In financial market models, agents use past prices to form idiosyncratic forecasts of future prices, and trade accordingly. New prices are created. Over time forecasting rules evolve, unprofitable ones are replaced by speculative ones, and the population of agents co-evolves to one another. In more traditional, non-financial market models prices are similarly emergent from the local interactions of self-interested agent. Interesting dynamics can result from the evolution of preferences, or shocks to supply chains, inventories, or technology, to name but a few possibilities. This distributed price formation process seems much closer to what happens in real world markets than the metaphor of the Walrasian auctioneer.

There is another notion of complexity at work here, one more in the spirit of the other papers in this Feature. Agent computing lies at the heart of the complex adaptive systems approach to complexity in economics. For such agent models are capable of producing perpetual dynamics at the agent level that yield coherent macrostructure that is at least very difficult to analyse analytically. From the perspective of agent modelling, the focus of analytical models on fixed points is mere mathematical expediency. For when one places purposive agents in economic environments of significant complexity, rarely do they stumble into any kind of equilibrium configuration. Rather, they engage in a kind of perpetual co-evolution to one another's strategies and there emerges more or less stable 'ecologies' of strategies. These ecologies may display quasi-stationary states but eventually such configurations are 'tipped' into other arrangements, i.e., the macro-equilibria are punctuated by periods of rapid transition.

In what follows we compare the computational complexity of the Walrasian model of exchange with a radically more decentralised one. In particular, Section 2 recapitulates recent results on the computational complexity of algorithms for computing Brouwer and Kakutani fixed points. It turns out that these algorithms fall into a complexity class that makes them among the hardest problems in all of computer science. Section 3 first describes the analytical structure of a general model of local exchange between agents, demonstrates that it can produce equilibrium allocations that are Pareto optimal, although not in the core and path dependent, and that the rate of convergence is geometric. Then, in Section 4 the complexity of this exchange process is investigated both analytically and compu-

¹ A more conventional title for this paper would have been 'Markets and Price Formation' but I feel there is a subtle bias in this terminology. For 'price formation' seems to presuppose that a single price characterises most exchange activity, and this is precisely what I wish to deny in the distributed, decentralised view of markets presented here.

tationally. It is demonstrated that the number of interactions required to produce an epsilon approximation of general equilibrium is polynomial in the number of agents and commodities. Section 5 summarises these results.

2. The Complexity of Walrasian Exchange

Fixed-point theorems were ostensibly introduced into economic theory by von Neuman (1945–6) in his work on the input-output model that bears his name. Since then, many domains of economic theory have come to depend on fixed point theorems to prove the existence of equilibria, notably general equilibrium theory but also Nash equilibria in game theory. For a recent perspective on this, see Geanakoplos (2003).

Now, the existence of an equilibrium is not the same as its achievement. That is, the achieving equilibrium requires a *mechanism* for converging to (an epsilon approximation of) a fixed point in a finite length of time, using a bounded (presumably small) amount of resources. Without such a mechanism there is little reason to believe that a fixed point would ever be observed.

The ‘tatonnement’ process of price adjustment is a *mechanism* for producing Walrasian general equilibrium. However, it is not a particularly realistic mechanism – among its several unreasonable requirements are that (1) agents truthfully reveal their preferences, (2) no trading takes place before the market-clearing price vector is announced, and (3) all agents trade at exactly the same prices. But the Walrasian mechanism has many nice properties: (1) it is *determinate* in the sense that the final prices and allocations are completely determined from agent preferences and endowments, and (2) the agent behaviour it requires is very simple, involving nothing more than truthful reporting of demands at announced prices.

Unfortunately, the job of the Walrasian ‘auctioneer’, who must compute prices, is extremely hard. The lower bound for worst-case computation of Brouwer fixed points is exponential in the dimension of the problem (Hirsch *et al.*, 1989) – the dimension being the size of the commodity space in the Arrow-Debreu version of general economic equilibrium. Furthermore, it has recently been shown that the computational complexity of Brouwer and Kakutani fixed points are closely related to the complexity of the parity argument, the connection between the two being Sperner’s lemma (Papadimitriou, 1994).² The constructive problem arising from the application of Sperner’s Lemma to the Brouwer and Kakutani fixed points of the Walrasian equilibrium model is that there are no polynomial time algorithms for the general case with nonlinear utility functions.³ Polynomial time algorithms

² The Brouwer fixed point $f(x^*) = x^*$ can be shown to involve the generalisation of Sperner’s Lemma for the N -dimensional simplex which guarantees the existence of x^* at the limit of the centre of the panchromatic simplices obtained by finer and finer triangulations, Papadimitriou (1994). See, Markose in this Feature, for further discussion on this and polynomial time or class-P algorithms.

³ More recent work of Papadimitriou and co-workers (Devanur *et al.*, 2002 and Deng *et al.*, 2002) pertain to linear utilities and demonstrate that polynomial algorithms exist for this restricted class of economies.

referred to as class-P problems are those that can be realistically solved by computers. These results can be summarised as

PROPOSITION 1: Arrow-Debreu equilibria are sufficiently difficult to compute that the Walrasian picture of market behaviour is simply not plausible.

There are at least two possible responses to this state of affairs. One is to simply dismiss theoretical complexity results, in the same way that exponential worst case complexity for linear programming does not vitiate use of the simplex algorithm for practical problems. Typical running times for particular general equilibrium codes have been estimated. Scarf (1973), for example, reports that the number of function evaluations required to equilibrate a computable general equilibrium (CGE) model via his algorithm scales like the size of the commodity space to the fourth power.⁴ While not as bad as exponential dependence, this result means that an economy with 1,000 commodities requires 10,000 times as many computations to equilibrate as compared with one with but 100 commodities. Such results seem unrealistic as a description of actual market behaviour.

A very different response is to argue that the Walrasian model, which has no empirical underpinnings (Hausman, 1992, p. 55), is not a reasonable picture of how an exchange economy works. For indeed, there are a variety of non-Walrasian exchange mechanisms that yield equilibrium allocations that are Pareto optimal. In particular, mechanisms that are radically more decentralised⁵ than the Walrasian one, with its single, uniform price vector, display greater fidelity to real economic processes. As Rust (1998) has written:

The reason why large scale computable general equilibrium problems are difficult for economists to solve is that they are using the wrong hardware and software. Economists should design their computations to mimic the real economy, using massively parallel computers and decentralised algorithms that allow competitive equilibria to arise as 'emergent computations'...[T]he most promising way for economists to avoid the computational burdens associated with solving realistic large scale general equilibrium models is to adopt an 'agent-based' modelling strategy where equilibrium prices and quantities emerge endogenously from the decentralised interactions of agents.

We will show that such decentralised exchange processes can have complexity properties that are better (less complexity) than the Walrasian process. If so, a further argument against Walras is a simple evolutionary one: if computation is costly then when two equally efficient (i.e., Pareto optimal) market mechanisms having significantly different computational complexity are competing, the one

⁴ The number of function evaluations does not depend on the number of agents, since the auctioneer uses only aggregate demand functions. However, this does not mean that the complexity of Walrasian equilibrium is independent of the number of agents. Rather, if one also accounts for how these demand functions are built up from the demands of individual agents then one gets that Walrasian equilibrium has complexity that is linear in the size of the population.

⁵ There is a long tradition here following the classic papers of Feldman (1973), Fisher (1972, 1989), Goldman and Starr (1982).

ultimately selected will be the least costly one, i.e., the one requiring the least number of computations – the market institution having the lower complexity.⁶

3. Decentralised Exchange Processes

This Section shows that a particular process of decentralised exchange in a population of agents having continuous, strictly convex preferences converges to an equilibrium. The bilateral exchange process is stable in the sense that starting from any initial conditions it always converges. Furthermore, convergence occurs at a geometric rate. Conditions under which such equilibria are Pareto optimal are given. In formulating these results a variety of previous work is synthesised.

Note the set of agents by $\mathcal{A} = \{1, \dots, A\}$ and the set of commodities by $\mathcal{N} = \{1, \dots, N\}$. Exchange occurs at a set of times $\mathcal{T} = \{1, 2, \dots, \tau\}$; elements of \mathcal{T} represent the indices of the sequence of physical times at which trade takes place. Each agent possesses an allocation $x^i(t) \in \mathbf{R}_+^N$ at each time $t \in \mathcal{T}$; $x^i(0)$ is agent i 's endowment. Each agent has continuous, strictly convex preferences, represented for agent i by utility function $U^i: \mathbf{R}_+^N \rightarrow \mathbf{R}$. Some number of agents, k , group to trade some number of goods at each period $t \in \mathcal{T}$. In general, multiple groups of agents can trade multiple goods at a particular time, but no agent can be a member of more than one trade group at a particular time. Call π^t the set of all agent groups that engage in trade at time $t \in \mathcal{T}$. A trade *history*, $\Pi = \{\pi^1, \pi^2, \dots, \pi^T\}$, gives the agent groups that trade particular goods at particular times; Π may be either exogenous or endogenous to the exchange process. Overall, the exchange process is given by the history-parameterised mapping $T_\Pi: \mathbf{R}_+^{AN} \rightarrow \mathbf{R}_+^{AN}$, that is

$$x(t+1) = T_\Pi[x(t)]. \quad (1)$$

Exchange between agents is required to be individually rational, that is, for each agent group, $\gamma \in \pi^t$, $U^k[x^k(t+1)] \geq U^k[x^k(t)]$ for all $k \in \gamma$, and or $U^k[x^k(t+1)] > U^k[x^k(t)]$ for some $k \in \gamma$. Exchange does not alter the total quantity of commodities. Define exchange to be *feasible* in a population of agents if there exists an agent group such that an individually rational exchange between the agents is possible. The agent population is in economic equilibrium when exchange is not feasible.

3.1. Existence of k -lateral Exchange Equilibria

Existence of k -lateral exchange equilibrium is easily demonstrated through construction of a Lyapunov function $V: \mathbf{R}_+^{AN} \rightarrow \mathbf{R}$ for the exchange process (Uzawa, 1962):

$$V[x(t)] \equiv \sum_{i \in \mathcal{A}} U^i[x^i(t)].$$

⁶ It is worth pointing out that Walras' notion of 'groping' for market-clearing prices more closely resembles the kind of decentralised exchange processes described below, so it is the Arrow-Debreu formalism that is being argued against here.

PROPOSITION 2: *k*-lateral exchange equilibria exist since

- (i) $V[x(t)]$ increases monotonically as long as trade takes place,
- (ii) the allocation path, $x[t, x(0)]$, is bounded, thus $V[x(t)]$ is bounded above;
- (iii) therefore

$$\lim_{t \rightarrow \infty} V\{x[t, x(0)]\} = V^*,$$

and $x[t, x(0)]$ has a subsequence converging on x^* such that $V(x^*) = V^*$.

Note that this result does not depend on any particular bargaining algorithm studied in for example Lengwiler (1994). As long as each trade is individually-rational then $V(t)$ is an increasing function.

3.2. Rate of Convergence of *k*-lateral Exchange Processes

The existence of equilibria is of little practical value if such equilibria are difficult to achieve, such as when prices cycle in the Walrasian adjustment process as first described by Scarf (1960). In the case of *k*-lateral exchange no such difficulties are encountered. Since each exchange makes at least one agent strictly better off, i.e., $U^i[x^i(t+1)] > U^i[x^i(t)]$ for some *i* and each *t*, it is also true that $V[x(t+1)] > V[x(t)]$ for all *t*. Define a convergence parameter, β , as

$$\limsup_{t \rightarrow T} \frac{V^* - V[x(t+1)]}{V^* - V[x(t)]} \equiv \beta.$$

As *t* increases, the above yields a non-increasing sequence of quotients, each less than one, implying that $\beta < 1$. Therefore,

$$\begin{aligned} V^* - V[x(t+1)] &\leq \beta\{V^* - V[x(t)]\} \\ &\leq \beta^{t+1}\{V^* - V[x(0)]\}. \end{aligned} \tag{2}$$

We have demonstrated:

PROPOSITION 3: *The rate of convergence of k-lateral exchange processes is geometric.*

This result will serve as the basis for the results on the computational complexity of decentralised in Section 4 below.

3.3. Stability of *k*-lateral Exchange Processes

It is usual in economic theory to talk about the stability of an exchange process, not of the allocations resulting from such a process. For example, in the context of the Walrasian model an auctioneer's rule is called *globally stable* if a price path approaches an equilibrium for all initial price vectors (Arrow and Hahn, 1971, p. 271). In the case of *k*-lateral exchange the existence of a Lyapunov function for the dynamics guarantees that every initial allocation will result in an equilibrium allocation.

PROPOSITION 4: *The k -lateral exchange process is globally stable.*

The more usual notion of stability – that of a perturbed dynamic system returning to equilibrium – never obtains for economic equilibria, since displacements in allocations that yield utility increases for some at the expense of others do not have individually rational paths back to the original equilibrium.

3.4. *Optimality of k -lateral Exchange Equilibria*

After two agents engage in welfare-improving trade their marginal rates of substitution (*MRSs*) in the exchanged commodities will be closer together than before trade. When the agents trade all the way to the contract curve their post-exchange *MRSs* will be identical and the allocations are optimal. What are sufficient conditions such that k -lateral optimal allocations throughout a population are equivalent to Pareto optimal allocations?

There are a variety answers to this question in the literature. The first was given by Rader (1968) and amounts simply to the requirement that one agent must have positive quantities of all commodities. This result is usually interpreted as the importance of having middlemen, market makers, and other types of agents who facilitate trade. The second answer was given by Feldman (1973). He showed that as long as all agents possessed some non-zero amount of a particular commodity then pairwise optimality implied Pareto optimality. Such a commodity is commonly interpreted as money. These results can be summarised in the following proposition, which is a kind of welfare theorem of decentralised exchange:

PROPOSITION 5 (*First welfare theorem for decentralised exchange*): *k -lateral exchange equilibria are Pareto optimal if either*

- (i) $\exists i$ s.t. $x_j^i > 0$ $\forall j$;
- (ii) $\exists j$ s.t. $x_j^i > 0$ $\forall i$.

The first condition may be interpreted as the existence of a middleman who holds all goods, while the second is the existence of money. These older results apply primarily to a population in which all agents can interact with one another, i.e., a perfectly mixed population. More recently, Bell (1997) gives analogous results for agents who interact over fixed networks.

This result is directly analogous to the first welfare theorem of neoclassical economics. But note that the distributed, decentralised character of the ‘invisible hand’ is manifest here. Indeed, the fact that the Smithian ‘hand’ is ‘invisible’ means that this version of it is much more in keeping with its intuitive meaning (Nozick, 1994; Rothschild, 1994).

The second welfare theorem states that any Pareto optimal allocation is a Walrasian equilibrium from some endowments, and is usually taken to mean that a social planner/society can select the allocation it wishes to achieve and then use tax and related regulatory policy to alter endowments such that subsequent market processes achieve the allocation in question. We have demonstrated above that the

job of such a social planner would be very hard indeed, and here we ask whether there might exist a computationally more credible version of the second welfare theorem.

First, note that the second welfare theorem invites the interpretation that endowments can be modified. In addition to preferences and endowments, in decentralised trade models the history of interaction determines final prices and allocations. Therefore, if we could somehow specify or alter the trade history the equilibrium outcome could be modified.

PROPOSITION 6: (*Second welfare theorem for decentralised exchange*): *Any Pareto optimal and individually rational allocation can be achieved via some decentralised exchange process.*

Results of this type are often found under the title of ‘accessibility of Pareto optima’ and are an active topic of research, see Bottazzi (1994). Note that Proposition 6 is more like a true converse of Proposition 5 than in the Walrasian model.

3.5. *Non-core Character of k-lateral Exchange Processes*

Although k -lateral exchange allocations are Pareto optimal, it is easy to see that they are not in the core (from initial endowments). Start two agents out with identical preferences and endowments and let them trade to equilibrium but with distinct interaction sequences. It would be mere coincidence if they ended up with identical allocations and thus a non-core allocation has been generated.⁷ Feldman (1973) first pointed this out for the bilateral exchange case.

PROPOSITION 7: *Allocations resulting from k-lateral exchange processes are not in the core*

It is ‘wealth effects’ which are the subject of the next subsection.

3.6. *Wealth Effect in k-lateral Exchange Processes*

While Walrasian exchange has no effect on the wealth of individual agents – that is, Walras’ law holds – in distributed exchange environments some agents gain wealth while others lose it. This is so because exchange at non-equilibrium prices alters agent wealth with respect to the equilibrium price. While the overall amount of wealth in the agent population is constant (at final market prices), the general effect of k -lateral exchange is to disperse wealth.

PROPOSITION 8: *k-lateral exchange processes disperse wealth if the following condition holds:*

$$\sum_{i=1}^A \{ \Delta w^i(t) [2w^i(0) + \Delta w^i(t)] \} \geq 0.$$

To establish this, compute the change in the variance as

⁷ Core allocations always have the equal treatment property, Green (1972).

$$\begin{aligned}
\text{var}[w(t)] &\geq \text{var}[w(0)] \\
\sum_{i=1}^A [px^i(t)]^2 &\geq \sum_{i=1}^A [px^i(0)]^2 \\
\sum_{i=1}^A \{ [px^i(t)]^2 - [px^i(0)]^2 \} &\geq 0 \\
\sum_{i=1}^A \{ [px^i(t) - px^i(0)] [px^i(t) + px^i(0)] \} &\geq 0 \\
\sum_{i=1}^A \{ \Delta w^i(t) [2w^i(0) + \Delta w^i(t)] \} &\geq 0.
\end{aligned}$$

Thus, under certain conditions the distribution of wealth can be expected to increase as a result of decentralised exchange.

3.7. Path Dependence of k -lateral Exchange Processes

Each distinct trade history will in general produce a distinct equilibrium. Since there are a combinatorially huge number of histories, there will exist vast numbers of k -lateral exchange equilibria. The vast majority of these are not accessible via a Walrasian mechanism (unless one rearranges endowments) since they do not satisfy the equal treatment property.

PROPOSITION 9: *Equilibrium allocations and prices depend on the history of exchange.*

The simple example mentioned above – of agents with identical endowments and preferences – well illustrates this path dependence.

4. Complexity of Decentralised Exchange

In this Section the complexity of bilateral exchange models is investigated. We present results on the number of computations required to achieve bilateral exchange equilibria as a function of the number of agents and the number of commodities. First some formal results are developed. Then computational results which support the formal analysis are given, for economies as large as a million agents and 20,000 commodities per agent.

4.1. Analytical Results

From the basic iteration (1) above, together with the fact that $\beta < 1$, we know that convergence is geometrically fast. It remains to figure out how β depends on the number of agents and the size of the allocation space, AN . Since this is a conservative system the operator T always has a unit eigenvalue, and so the rate of convergence, β , is controlled by the sub-dominant eigenvalue. For particular exchange processes it is possible to compute β explicitly. In general it is possible to

place an upper bound on the number of interactions required to equilibrate a market by noting that each application of T requires no more than $O[(AN)^2]$ operations. Define

$$\varepsilon \equiv \frac{V^* - V[x(t)]}{V^* - V[x(0)]} \leq \beta^\tau$$

and solve for τ the number of time steps required to produce an ε approximation of equilibrium. This quantity is clearly bounded from above by

$$\tau \leq \frac{\ln(1/\varepsilon)}{\ln(1/\beta)}.$$

This leads naturally to

PROPOSITION 10: *The number of interactions is bounded from above by $A^2N^2\tau(\varepsilon, \beta)$, therefore the computational complexity of k -lateral exchange is P .*

Furthermore, in the case of bilateral exchange it is possible to develop a sharper result.

PROPOSITION 11: *The computational complexity of bilateral exchange is bounded from above by $AN^2\tau(\varepsilon, \beta)$.*

There are at least three ways to understand these results heuristically. First, because the dimension of the allocation space is AN , each interaction shrinks the set of feasible bilateral exchanges by

$$\beta_0^{1/AN},$$

where β_0 is some constant. But this is not the whole story, for it is also true that as N gets larger, the number of MRS s increases linearly. Thus, the number of interactions required to converge a norm of the vector of MRS s to within ε of equilibrium scales like ε^N . Therefore, calling \bar{I} the maximum number of bilateral interactions necessary to reach an ε approximation of equilibrium, we have

$$\bar{I} \leq \frac{\ln(1/\varepsilon_0^N)}{\ln(1/\beta_0^{1/AN})} = \frac{N \ln(1/\varepsilon_0)}{1/AN \ln(1/\beta_0)} = AN^2 \frac{\ln(1/\varepsilon_0)}{\ln(1/\beta_0)}.$$

Overall, this upper bound on the requisite number of interactions is proportional to the number of agents and the number of commodities squared.

The second intuitive rationale for these results relates to the dependence on A . Imagine an economy composed of very large numbers of both agents and commodities, and consider two experiments. First, randomly divide the population into two equal-sized groups of agents, each of whom has preferences over the entire commodity space. Now equilibrate each one via bilateral exchange with agents in its own group. Each of these sub-economies converges to very similar prices. In fact, if the overall economy is large enough then the two sub-economies

converge to exactly the same price, in which all agents have the same MRSs. Thus, combining the two groups of agents subsequently there are very few further trades that can be arranged and the economy is quickly put in equilibrium.

The third thought experiment that conveys the general character of the result consists of converging the economy across all agents for two commodities and measuring the number of total interactions required – call it I_2 . Now, to this equilibrium configuration add 1 commodity and re-equilibrate. Intuitively, the number of interactions necessary should be proportional to $2I_2$, since the new commodity must be equilibrated with each of the original two. Next, add a fourth commodity, requiring interactions proportional to $3I_2$, and so on until the commodity space consists of all N commodities. Overall, the total number of interactions necessary for this process is

$$\sum_{i=2}^N (i-1)I_2 = I_2 \sum_{i=2}^N (i-1) = I_2 \sum_{i=1}^{N-1} i = I_2 \frac{N^2}{2}.$$

Thus, the quadratic dependence on N .

Example 1: Bilateral exchange on a circle

Albin and Foley (1990) studied exchange of 2 commodities in parallel among agents with homogeneous Cobb-Douglas preferences arranged in a circle. Consider such a ring, composed of A agents, an even number, indexed from 1. In any period all even numbered first agents trade with odd numbered agents whose index is less than theirs, e.g., agent no.4 trades with agent no.3. Then, even numbered agents trade with the odd indexed agents just above them. Overall, this double set of trades constitute one time period. Call $T_1(x)$ the algorithm by which even-numbered agents trade with those directly below them, and $T_2(x)$ the other exchange process. Then the overall exchange algorithm is $T(x) = T_2[T_1(x)]$. It can be shown that this leads to a tridiagonal set of equations, for which the eigenvalues can be calculated explicitly; the subdominant one is strictly less than unity.

4.2. Computational Results

In this Section results from a variety of computational experiments are described involving bilateral exchange at local Walrasian prices. Overall, some one trillion exchange transactions are summarised here. These results support the analytical results obtained above.

4.2.1. Dependence on the number of agents

How many agent-agent interactions are required to produce a bilateral exchange equilibrium?⁸ This depends on many things, including how good an approximation to equilibrium we wish to compute. But the general character of the

⁸ We will use the number of interactions as a surrogate for the complexity of the exchange process. This is reasonable since each interaction involves a fixed number of computations.

dependence of the number of interactions on the number of agents will not be sensitive to the accuracy of the approximation.

Example 2: Heterogeneous Cobb-Douglas agents

Figure 1 gives the number of interactions necessary to equilibrate a heterogeneous population of agents having Cobb-Douglas preferences distributed uniformly over $(0, 1)$. Endowments are distributed uniformly over $[50, 150]$. Agents are paired at random, truthfully report their preferences, and trade directly to the contract curve. All exchange is terminated once the largest standard deviation in the $\ln(\text{MRS})$ distribution falls below 10^{-2} . Results are shown for three distinct sizes of the commodity space, $N = 2, 10$ and 50 , varying the number of agents from 10 to 1,000,000.

Note that the effect of increasing the number of commodities is merely to increase the number of interactions necessary for equilibration, but does not change the nature of the dependence on the number of agents. Each line in the figure has a slope of 1.000, meaning that as the number of agents increases the number of interactions required to produce equilibrium increases linearly. If 10^3 bilateral interactions are necessary to equilibrate a population of 10^2 agents then 10^6 agents require 10^7 interactions, and so on. The number of interaction/agent is independent of the size of the population.

Similar results obtain for different values of the termination criteria. The overall effect of decreasing ε is to require more agent-agent interactions, as shown in Figure 2 for $N = 2$.




Fig. 1. *Number of Interactions Required for Market Convergence as a Function of the Number of Agents, A, parameterised by the number of commodities, N; termination occurs once $|\text{variance} [\ln(\text{MRS})]|^\infty \leq \varepsilon = 0.01$*




Fig. 2. *Number of Interactions Required for Convergence as a Function of the Number of Agents, A , parameterised by the termination criterion ε , such that termination occurs once $|\text{variance} [\ln(\text{MRS})]|^\infty \leq \varepsilon$; 2 commodities*

For other specifications of preferences (e.g., CES) and different interaction topologies (e.g., parallel instead of serial) we have obtained results identical in character, i.e., the number of interactions required to produce bilateral exchange equilibria is linear in the population size.

This result has important implications. Imagine if it were not true, but rather that the number of interactions/agent *increased* as the total number of agents *increased*. Then, as each new agent were added to the society the economic complexity for each extant agent would grow, independent of whether or not any particular agent even interacted with the new agent. This seems unreasonable. Rather, bilateral trade produces a kind of *social computer* which endogenously decentralises economic computations.

4.2.2. *Dependence on the Number of Commodities*

The dependence of the number of interactions on the number of commodities is similar. In the example below we find that as the commodity space, N , increases the number of bilateral interactions required to produce equilibrium increases in proportion to N^2 .

Example 2 (continued):

We instantiate various populations of Cobb-Douglas agents, as above, having heterogeneous preferences and endowments, pair them randomly, and track the number of interactions necessary to produce bilateral exchange equilibria having a

variance of no more than 10^{-2} in the final $\ln(\text{MRS})$ distribution, all as a function of the number of commodities (N from 2 to 20,000). The results are shown in Figure 3.

The slope of each curve in Figure 3 is 2.000, meaning that the required number of interactions scales like N^2 . The effect of increasing the number of agents, holding the number of commodities constant, is merely to increase the number of interactions required for equilibrium.

The effect of tightening approximation is to require additional interactions, as shown in Figure 4.

An interesting open question is ‘Does there exist an exchange process for producing Pareto optimal allocations that has complexity linear in the number of commodities?’ From the computational evidence above it would appear that the answer is ‘no’ for bilateral exchange.

5. Summary and Conclusions

It has been argued that the Walrasian model of exchange is problematical on a variety of grounds. Notably, recent results on the computational complexity of Brouwer and Kakutani fixed points suggest that real markets cannot possibly operate according to the Walrasian model. A decentralised exchange model has been offered as an alternative to the Walrasian picture. In particular, k -lateral exchange equilibria have much better computational complexity than do Walrasian equilibria. Differences between the models are summarised in Table 1.




Fig. 3. *Number of Interactions Required for Convergence as a Function of the Number of Commodities, N , parameterised by the number of agents, A ; termination occurs once $||\text{variance}[\ln(\text{MRS})]||^\infty \leq \varepsilon = 0.01$.*

Interactions

$A = 100$




Fig. 4. *Number of Interactions Required for Convergence as a Function of the Number of Commodities, N , parameterised by the termination criterion ϵ , such that termination occurs once $\|\text{variance} [\ln(\text{MRS})]\|^\infty \leq \epsilon$; 100 agents*

As described above, Scarf has observed that for general CGE models the number of computations scales like N^4 . Call $k_b N^2$ the number of computations required to equilibrate a N commodity economy via bilateral exchange, while $k_W N^4$ is the corresponding number for Walrasian exchange. Clearly, if $k_b < k_W$ then the bilateral exchange process is always more efficient computationally than the Walrasian one. Consider the opposite case, $k_b > k_W$. Then for small numbers of commodities the Walrasian process requires fewer computations but as N grows

Table 1
Comparison of Walrasian and Decentralised Exchange Equilibria

	<i>Wabras-Arrow-Debreu</i>	<i>k-lateral Exchange</i>
<i>Price formation</i>	Global	Local
<i>Price determination</i>	OR problem (ostensibly solved by auctioneer)	DAI problem ⁹ ('solved' by market of agents)
<i>Existence of equilibrium</i>	Fixed point theorems	Lyapunov function
<i>Character of equilibrium</i>	Determinate (depends on preferences, endowments)	Indeterminate (depends also on interaction history)
<i>Welfare of equilibrium</i>	Pareto optimal	Pareto optimal
<i>Stability of equilibrium</i>	Ambiguous	Globally stable
<i>Dynamics</i>	One-shot (no trade out of equilibrium)	Path-dependent
<i>Wealth effect</i>	None (Walras' law holds)	Dispersive (Walras' law violated)
<i>Complexity</i>	Exponential (worst case) N^4 (average case)?	Polynomial (quadratic in A and N)

⁹ DAI stands for Distributed Artificial Intelligence.

the bilateral exchange process quickly becomes more efficient. There is some critical number of commodities, N_c such that

$$k_b N_c^2 = k_w N_c^4.$$

For $N > N_c = \sqrt{k_b/k_w}$, the bilateral exchange process is superior to the Walrasian one computationally. Note that even if the Walrasian algorithms are 100 times more efficient for small problems – i.e., $k_b/k_w = 100$ – bilateral exchange will be more efficient for N greater than 10.

Walrasian markets in their Arrow-Debreu conception are an ideal type, in the terminology of the philosophy of science, a caricature of reality that abstracts from many details of real markets in order to provide a home for our intuitions and a point of departure for deeper exploration of market processes. Unfortunately, the embodiment of this ideal type in CGE software, especially when utilised for policy purposes, institutionalises a series of propositions that more behaviourally realistic and decentralised models reveal to be false, namely, that markets do not disperse wealth, yield allocations that are determined solely by preferences and endowments and are not history-dependent. Luckily, the unreality of this ideal type is given away by its computational intractability.

In the end we advocate not the jettisoning of this useful abstraction but merely its circumspect use whenever focused on questions for which it has limited ability to adjudicate an appropriate answer, e.g., distributional issues, actual prices. But because policy-focused model deal always and everywhere with just these issues, a direct consequence of the results described above is to at least cast a pale on the utility of such analyses, if not to vitiate them altogether.

The Brookings Institution

References

- Albin, P. and Foley, D. K. (1990). 'Decentralized, dispersed exchange without an auctioneer: a simulation study', *Journal of Economic Behavior and Organization*, vol. 18(1), pp. 27–51.
- Arrow, K. J. and Hahn, F.H. (1971). *General Competitive Analysis*, New York: North-Holland.
- Arthur, W. B., Durlauf, S.N. and Lane, D.A. (1997). *The Economy as a Complex Evolving System*, vol. 2, Boston MA: Westview Press.
- Bell, A. M. (1997). 'Bilateral trading on a network: convergence and optimality results', Working paper No. 97-W08, Vanderbilt University.
- Bertsekas, D. and Tsitsiklis, J. (1989). *Parallel and Distributed Computation*, New York: Wiley.
- Bottazzi, J.-M. (1994). 'Accessibility of Pareto optima by Walrasian exchange processes', *Journal of Mathematical Economics*, vol. 23, pp. 585–603.
- Brewer, P. J. (1999). 'Decentralized computation procurement and computational robustness in a smart market', *Economic Theory*, vol. 13, pp. 41–92.
- Chen, S.-H. and Yeh, C.-H. (1997). 'Modeling speculators with genetic programming', in (P. Angeline, R. G. Reynolds, J. R. McDonnell and R. Eberhart, eds.), *Evolutionary Programming*, vol. 6, New York: Springer-Verlag.
- Cheng, J. and Wellman, M. (1998). 'The WALRAS algorithm: a convergent distributed implementation of general equilibrium outcomes', *Computational Economics*, vol. 12, pp. 1–24.
- Darley, V., Outkin, A. Plate, T. and Gao, F. 2001, 'Learning, evolution and tick size effects in a simulation of the NASDAQ stock market', in *Proceedings of the 5th World Multi-Conference on Systemics, Cybernetics and Informatics*, Orlando, Florida.
- Delli Gatti, D., Gallegati, M. and Kirman, A. P. (eds.) (2000) *Interaction and Market Structure. Essays on Heterogeneity in Economics*, Lecture Notes in Economics and Mathematical Systems, vol. 484, Berlin: Springer.

- Deng, X., Papadimitriou, C. and Safra, S. (2002). 'On the complexity of equilibria', *Proceedings of the 34th Annual Association for Computing Machinery (ACM) Symposium on Theory of Computing*, NY: ACM Press.
- Devanar, N. R., Papadimitriou, C. H., Saberi, A. and Vazirani, V. V. (2002). 'Market equilibrium via a primal-dual-type algorithm', 43rd Annual IEEE Symposium on Foundations of Computer Science (FOC '02), Vancouver, B.C.
- Epstein, J. M. and Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*, Cambridge, MA: MIT Press.
- Feldman, A. (1973). 'Bilateral trading processes, pairwise optimality, and Pareto optimality', *Review of Economic Studies*, vol. 40(4), pp. 463–73.
- Fisher, F. M. (1972). 'On price adjustment without an auctioneer', *Review of Economic Studies*, vol. 39, pp. 1–15.
- Fisher, F. M. (1989). 'It takes ϵ^* to tango: trading coalitions with fixed prices', *Review of Economic Studies*, vol. 56, pp. 391–404.
- Geanakopolos, J. (2003). 'Nash and Walras from Brouwer', *Economic Theory*, vol. 21, pp. 585–603.
- Goldman, S. M. and Starr, R.M. (1982). 'Pairwise, ϵ -wise, and Pareto optimalities', *Econometrica*, vol. 50(3), pp. 593–606.
- Green, J. R. (1972). 'On the inequitable nature of core allocations', *Journal of Economic Theory*, vol. 4, pp. 132–43.
- Hausman, D. M. (1992). *The Inexact and Separate Science of Economics*, New York: Cambridge University Press.
- Hirsch, M. D., Papadimitriou, C. H. and Vavasis, S. A. (1989). 'Exponential lower bounds for finding Brouwer fixed points', *Journal of Complexity*, vol. 5, pp. 379–416.
- LeBaron, B. (2001a). 'Financial market efficiency in a coevolutionary environment', In *Agents 2000: The Simulation of Social Agents: Architectures and Institutions*, Chicago, Ill: Argonne National Laboratory, available online.
- LeBaron, B. (2001b). 'A builder's guide to agent-based financial markets', *Quantitative Finance*, vol. 1, pp. 254–61.
- LeBaron, B. (2001c). 'Evolution and time horizons in an agent-based stock market', *Macroeconomic Dynamics*, vol. 5, pp. 225–54.
- Lengwiler, Y. (1994). 'Bilateral economies: a game-theoretic analysis of general economic equilibrium', doctoral dissertation (no. 1554), Hochschule St. Gallen, Switzerland.
- Levy, H., Levy, M. and Solomon, S. (2000). *Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena*, New York: Academic Press.
- Lux, T. (1998). 'The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions', *Journal of Economic Behavior and Organization*, vol. 33, pp. 143–65.
- von Neumann, J. (1945–6). 'A model of general economic equilibrium', *Review of Economic Studies*, vol. 13(33), pp. 1–9, translation of 'Über ein ökonomisches Gleichungssystem und eine Verallgemeinerung des Brouwerschen Fixpunktsatzes', *Ergebnisse eines Mathematischen Kolloquiums*, vol. 8, pp. 73–83.
- Nozick, R. (1994). 'Invisible hand explanations', *American Economic Review Papers and Proceedings*, vol. 84(2), pp. 314–8.
- Papadimitriou, C. H. (1994). 'On the complexity of the parity argument and other inefficient proofs of existence', *Journal of Computer and System Sciences*, vol. 48, pp. 498–532.
- Rader, T. (1968). 'Pairwise optimality and non-competitive behavior', in (J. P. Quirk and A. M. Zarley, eds.), *Papers in Quantitative Economics*, Lawrence KS: University Press of Kansas.
- Rothschild, E. (1994). 'Adam Smith and the invisible hand', *American Economic Review Papers and Proceedings*, vol. 84(2), pp. 319–22.
- Rust, J. (1998). 'Dealing with the complexity of economic calculations', in (S. Durlauf and J. Traub, eds.), *Limits to Knowledge in Economics*, Boston MA: Addison-Wesley.
- Scarf, H. E. (1960). 'Some examples of global instability of the competitive equilibrium', *International Economic Review*, vol. 1, pp. 157–72.
- Scarf, H. (with T. Hansen) (1973). *Computation of Economic Equilibrium*, New Haven, Connecticut: Yale University Press.
- Uzawa, H. (1962). 'On the stability of Edgeworth's barter process', *International Economic Review*, vol. 3(2), pp. 218–32.
- Zawadowski, A. G., Karádi, R. and Kertész, J. (2002). 'Price drops, fluctuations, and correlation in a multi-agent model of stock markets', Working paper arXiv:cond-mat/0204574v1.