

MAXIMUM ENTROPY EXCHANGE EQUILIBRIUM

Duncan K. Foley

Department of Economics, Barnard College, Columbia

University, New York, NY 10027 USA*

1. Statistical and Walrasian concepts of economic equilibrium

The idea of treating the state of an economy as an equilibrium analogous to the equilibrium of a mechanical system has had enormous influence on the development of economic theory (see Philip Mirowski (1990)). The purpose of this paper is to explore the consequences of applying the methods of statistical mechanics to the standard economic model of pure exchange when the number of agents of each type characterized by tastes and endowment is large. The economic analogue of statistical mechanical equilibrium in the pure exchange model is the *maximum entropy exchange equilibrium*, the Pareto-efficient allocation that can be realized in the largest number of ways through permutation of agents of the same type under the constraint that all exchanges are voluntary.

The maximum entropy exchange equilibrium shares with the concept of Walrasian competitive equilibrium the property of Pareto-efficiency: both concepts represent allocations in which there are no further opportunities for mutually beneficial trade because of the effects of competition. The maximum entropy exchange equilibrium of an exchange economy differs from the Walrasian competitive equilibria of the same economy in allowing for the effects of bilateral trades at disequilibrium prices, whereas Walrasian competitive equilibrium explicitly rules out trading at disequilibrium

* This paper grew out of conversations with Chris Grandy, Perry Mehrling, and Jed Harris. I am particularly indebted to Grandy for making available parts of Jaynes' manuscript.

prices through the fictional device of an auctioneer who centralizes the informational aspects of market exchange. At the maximum entropy exchange equilibrium agents of the same type have different final consumption bundles, depending on their individual gains from the trading process, while at a Walrasian competitive equilibrium all agents of a given type have the same final consumption bundle. Thus the maximum entropy exchange equilibrium provides a tool for studying inequality that is endogenously generated by markets.

A unique maximum entropy exchange equilibrium exists when preferences are strictly convex.

Whereas the maximum entropy exchange equilibrium is the most probable Pareto-efficient, individually utility improving state of an exchange economy, Walrasian competitive equilibrium turns out to have a vanishingly small probability. Thus if we observe Walrasian competitive equilibrium, there must be very powerful forces sustaining it as an allocation.

2. The pure exchange economy

There are $m < \infty$ commodities, so that commodity bundles can be represented by points in R^m : $x = (x_1, \dots, x_m)$. There are r types of agents, denoted by superscript k , each represented by a continuous, strictly increasing, strictly concave utility indicator $u^k: X^k \rightarrow [-\infty, \infty)$, and an endowment $\omega^k \in R^n$. (The utility indicator takes the value $-\infty$ at infeasible consumption plans.) The number of agents of type k is n^k , $n = \sum_{k=1}^r n^k$ is the total number of agents in the economy, and $w^k = \frac{n^k}{n}$ is the weight of type k in the economy. The agents are distinguished by superscript i , with the convention that agents of type 1 come first, type 2 next, and so on. An allocation $x = (x^1, x^2, \dots, x^{n^1}, x^{n^1+1}, x^{n^1+2}, \dots, x^{n^1+n^2}, \dots, x^n) \in R^{mn}$ assigns a commodity bundle to each agent. Thus $x = (\omega^1, \dots, \omega^1, \omega^2, \dots, \omega^2, \dots, \omega^n, \dots, \omega^n)$ with each endowment

repeated n^k times, is the original endowment. $\mathbf{u}[\mathbf{x}] = (u^1[x^1], \dots, u^n[x^n])$ is the vector of utilities achieved by the allocation \mathbf{x} . It is convenient to write the resources required per agent for an allocation \mathbf{x} as $\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}^i$. Thus

$\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}^i$ is the average endowment of the whole economy.

Three important properties of allocations are *feasibility*, *arbitrage*, and *individual improvement*.¹

Feasibility requires that the allocation assign all the available commodities to some agent.

Definition 1 (feasibility): An allocation \mathbf{x} is

feasible if $\bar{\mathbf{x}} = \bar{\mathbf{x}}$.

Arbitrage requires that there are no further mutually improving trades. A net trade \mathbf{y} between agents i and j leaves agent i with $\mathbf{x}^i + \mathbf{y}$ and agent j with $\mathbf{x}^j - \mathbf{y}$.

Definition 2 (arbitrage): An allocation \mathbf{x} is *arbitrated* if for any two agents i, j , and any nonzero vector \mathbf{y} , either $u^i[\mathbf{x}^i + \mathbf{y}] < u^i[\mathbf{x}^i]$ or $u^j[\mathbf{x}^j + \mathbf{y}] < u^j[\mathbf{x}^j]$, or both.

Corollary 2.1: If an allocation \mathbf{x} is arbitrated, then under the hypothesis of concave utility indicators, there exists a nonempty set of *supporting price systems* to \mathbf{x} , $P[\mathbf{x}]$

$$= \sum_{i=1}^n \frac{\partial u^i[\mathbf{x}^i]}{|\partial u^i[\mathbf{x}^i]|} \neq \emptyset.$$

Proof: The sum of the sets of net utility-increasing trades for all agents is a convex set whose interior does not contain the

¹ $\mathbf{x} \geq \mathbf{x}'$ means $x_j \geq x'_j$ for all j ; $\mathbf{x} \geq \mathbf{x}'$ means $\mathbf{x} \geq \mathbf{x}'$ and $\mathbf{x} \neq \mathbf{x}'$. If $u: \mathbb{R}^m \rightarrow [-\infty, \infty)$, $\partial u[\mathbf{x}]$ is the subgradient of u at the point \mathbf{x} . For price systems $\mathbf{p} \geq \mathbf{0}$, the norm $|\mathbf{p}| = \sum_{j=1}^m p_j$

. Unless otherwise noted, price systems are assumed to have $|\mathbf{p}| = 1$.

origin. Therefore a price system exists separating the origin from the sum set, which is the required common element in the subgradients.

Corollary 2.2: An allocation is feasible and arbitrated if and only if it is Pareto-efficient.

Definition 3 (individual improvement): An allocation \mathbf{x} is *individually improving* if $\mathbf{u}[\mathbf{x}] \geq \mathbf{u}[\mathbf{0}]$.

Voluntary bilateral trade can achieve only feasible and individually improving allocations. There is no economic motivation for agents to alter an arbitrated allocation through exchange. Thus feasible, arbitrated and individually improving allocations are the equilibria of voluntary trading systems.

Definition 4 (micro-exchange equilibrium): A feasible, arbitrated and individually improving allocation \mathbf{x} is a *micro-exchange equilibrium*.

Since a micro-exchange equilibrium \mathbf{x} is arbitrated, by Corollary 2.1 there is at least one system of equilibrium prices naturally associated with it, $\mathbf{p}[\mathbf{x}] \in P[\mathbf{x}]$, a common element in the subgradients of the utility functions at the equilibrium allocation.

Walrasian competitive equilibrium is a micro-exchange equilibrium that can be reached by bilateral trades made at the eventual equilibrium prices. Thus Walrasian competitive equilibrium prices give the same value to each agent's final consumption bundle and endowment.

Definition 5: (Walrasian competitive equilibrium): A micro-exchange equilibrium \mathbf{x} is a Walrasian competitive equilibrium if $\mathbf{p}[\mathbf{x}]\mathbf{x}^i = \mathbf{p}[\mathbf{x}]\mathbf{0}^i$ for all i .

3. Maximum entropy exchange equilibrium

There is in general a large number of micro-exchange equilibria. Since there is no economically meaningful

distinction among agents of the same type, two micro-exchange equilibria that differ only in the interchange of consumption bundles of agents of the same types are economically identical. The micro-exchange equilibria are thus naturally grouped into equivalence classes, *exchange equilibria*, each of which consists of micro-exchange equilibria that can be transformed into each other through the permutation of the consumptions bundles of identical agents.

If we were to observe a pure exchange economy, there would be limits to the fineness with which we could distinguish different utilities. In other words there would be some small *grain* $\epsilon > 0$ such that agents of the same type whose utilities differed by amounts smaller than ϵ would be observationally indistinguishable. Thus it is reasonable to divide the utility line for agents of type k into a large number U^k of segments of length ϵ , and to regard agents of type k whose utilities lie in the same segment as having the same utility. Since exchange equilibria are individually improving, we can start these segments at $u^k[\epsilon^k]$, the minimum utility any agent will bargain to voluntarily. Since the economy's endowment is finite, the maximum utility of each type of agent is finite, and we can choose U^k large enough to encompass all the relevant possibilities.

Given a price system p and utility level u , let $x^k[u;p]$ be the Hicksian demand system for type k , which, under the hypothesis of strict concavity of the utility indicators, is unique, and gives the consumption bundle for an agent of type k facing the price system p at utility level u .

Definition 6 (Hicksian demand functions):

Given a price system p and a utility level u ,

$$x^k[u;p] = \arg \min_{x \in \mathbb{R}^n} px \text{ subject to } u^k[x] \geq u.$$

When utility indicators are continuous, strictly increasing, and strictly concave, the Hicksian demand functions are unique (when they exist) and continuous as functions of price.

Since agents of the same type are observationally indistinguishable, two micro-exchange equilibria that differ only in the interchange of two agents of the same type are economically equivalent.

To make this equivalence transparent, we can calculate for any micro-exchange equilibrium \mathbf{x} the corresponding price system, $p[\mathbf{x}]$, and the number of agents of each type occupying each segment of the utility line in that equilibrium, $\{n_u^k[\mathbf{x}]\}$. Denote the proportion of the agents of type k with

utility u as $f_u^k = \frac{n_u^k}{n^k}$ so that, since $\sum_{u=u^k[\mathbf{x}]}^{u^k} n_u^k = n^k$, $\sum_{u=u^k[\mathbf{x}]}^{u^k} f_u^k$

$$\sum_{k=1}^r w_k \sum_{u=u^k[\mathbf{x}]}^{u^k} x^k[u;p] f_u^k = \bar{w} \quad (1)$$

(Feasibility implies)

Micro-exchange equilibria that can be transformed into each other by permuting the consumption bundles of agents of the same type evidently have the same p and $\{n_u^k\}$.

Definition 7 (exchange equilibrium): An *exchange equilibrium* $(p, \{n_u^k\})$ is the set of micro-exchange equilibria \mathbf{x} with $p[\mathbf{x}] = p$ and $n_u^k[\mathbf{x}] = n_u^k$.

The negative of the logarithm of the number of ways that an exchange equilibrium $(p, \{n_u^k\})$ can be achieved through permutation of consumption bundles of agents of the same type divided by the total number of agents is the *entropy* of the exchange equilibrium.

Definition 8 (entropy): The *entropy* of an exchange equilibrium $(p, \{n_u^k\})$ is $-n^{-1} \ln W(p, \{n_u^k\})$ where

$$W(p, \{n_u^k\}) \equiv \frac{n^1!}{n_1^1! \dots n_u^1! \dots n_{U^1}^1!} \dots \frac{n^r!}{n_1^r! \dots n_u^r! \dots n_{U^r}^r!}$$

If the bilateral trading process has no systematic properties beyond feasibility, arbitrage, and individual improvement, every micro-exchange equilibrium will be equally probable. But exchange equilibria are not all equally probable, because some can be realized in a much larger number of micro-exchange equilibria than others. If we were to observe a particular exchange economy, we would be most likely to find it in the most probable exchange equilibrium, the exchange equilibrium with the highest entropy.

Definition 6 (maximum entropy exchange equilibrium): The *maximum entropy exchange equilibrium* is the exchange equilibrium that has the highest entropy.

Two interpretive stipulations are crucial to understanding the limits of the concept of maximum entropy exchange equilibrium. First, the number of agents of each type must be very large, so that statistical considerations become of decisive importance. Second, the exchange process must be a once and for all event not embedded in any larger stationary process. The exchange economy can be thought of as a village market, in which every agent loads her total produce (her endowment) onto a wagon and brings it to a central place, makes bilateral trades with other agents without knowledge of the other trades taking place until there are no further trading opportunities to be found, and takes her final holding back home in the wagon. In this scenario, it is essential that the agent have no idea whether her endowment is close to or far from the average endowment of the village, and no idea where the equilibrium prices will settle. She is guided in her trading only by her interest in avoiding exchanges that reduce her utility. Stronger restrictions on the trading process, such as limitations on the size or exchange ratios of bilateral trades, would have to be introduced as explicit further assumptions in the definition of equilibrium, and would alter the maximum entropy exchange equilibrium.

The maximum entropy exchange equilibrium is the state of a pure exchange economy we would be most likely to observe in the absence of any further constraints on the final allocation. But it is important to see that there are no economic forces tending to move the economy to the maximum entropy exchange equilibrium from any other exchange equilibrium. If the economy were to find a feasible, arbitrated, individually improving allocation different from the maximum entropy exchange equilibrium, it would still be in equilibrium since there would be no further motivation for voluntary trade. This point reflects an important limitation of the model of pure exchange as a theory of real economic relations. Pure exchange is a one-time atemporal relaxation into equilibrium with no ongoing dynamics, whereas time structures production and consumption in real economies and generates systematic dynamics.

a. Existence and uniqueness

The problems of finding the maximum entropy exchange equilibrium can be approached in two steps. First, maximize the entropy for each given price system p with $|p| = 1$. This problem is:

$$\begin{aligned}
 H[p] &= \max_{(p, \{n_u^k\})} n^{-1} \ln W[(p, \{n_u^k\})] \\
 \text{subject to } & \sum_{u=u^k[\square^k]}^{U^k} n_u^k = n^k \quad \text{for } k = 1, \dots, r \\
 & \sum_{k=1}^r \frac{n^k}{n} \sum_{u=u^k[\square^k]}^{U^k} x^k[u;p] \frac{n_u^k}{n^k} = \bar{x} \quad \text{(MAXENTP)}
 \end{aligned}$$

Since for a given grain and U^k the number of possible states is finite, the number of possible assignments of the finite number of agents to those states is also finite. If there are any assignments that satisfy the constraints, there

must be at least one that maximizes entropy. Assign $H[p] = -$ to those price systems for which there is no feasible assignment.

The next step is to maximize $H[.]$ over the possible price systems. The study of this problem is simplified in the case of concave utility indicators.

Theorem 1: Let $H[p]$ be the maximum entropy for an exchange economy at the relative price system p . If the utility indicators, $u^k[.]$, are concave, $H[.]$ is continuous and strictly concave on its proper domain.

Proof: Continuity follows from the continuity of the Hicksian demand functions. Let p, p' be two relative price systems, with $H[p], H[p'] > -$, and $\{f_u^k\}$ and $\{f_u^{k'}\}$ the corresponding entropy maximizing distributions over u . For $\alpha \in (0,1)$ the distribution $\{\alpha f_u^k + (1-\alpha)f_u^{k'}\}$ satisfies the conservation of agents condition. Give a fraction α of the agents of type k in state u $x^k[u;p]$ and a fraction $(1-\alpha)$ $x^k[u;p']$. In this allocation all the agents of type k in state u are achieving utility level u . Furthermore, this allocation is feasible (though not necessarily arbitrated, and hence not necessarily Pareto-efficient). Since the entropy function is strictly concave over distributions, this distribution has a higher entropy than the average of the entropy of the two original distributions. The possibility of moving to a Pareto-efficient allocation will save some resources, making possible an even higher entropy distribution, so $H[\alpha p + (1-\alpha)p'] > \alpha H[p] + (1-\alpha)H[p']$.

There are two cases. If $H[p] = -$ for all price systems p , then there exist no micro-exchange equilibria for the economy (and no individually improving Pareto-efficient allocations). If $H[p] > -$, for some p , then the problem of finding the maximum entropy exchange equilibrium consists of maximizing a strictly concave function over a compact convex set, a problem which has a unique solution.

Thus we can solve:

$$f_u^k [p; \lambda, \lambda] = \exp[-1 - \lambda^k] \exp[-\lambda x^k [u; p]]$$

Summing over the u for each k eliminates the terms involving λ^k :

$$f_u^k [p; \lambda] = \frac{\exp[-\lambda x^k [u; p]]}{\int_{u=u^k[\lambda^k]}^{U^k} \exp[-\lambda x^k [u; p]] du}$$

Here $\lambda [p] \in \mathbb{R}^m$ must be chosen to satisfy the m feasibility conditions:

$$\sum_{k=1}^r w^k \int_{u=u^k[\lambda^k]}^{U^k} f_u^k [p; \lambda] x^k [u; p] du = \bar{x}$$

We can then calculate the mean utility and mean consumption vector for each type at the maximum entropy distribution for the price system p , $\{f_u^k [p]\}$:

$$\bar{u}^k [p] = \frac{\int_{u=u^k[\lambda^k]}^{U^k} f_u^k [p] u du}{\int_{u=u^k[\lambda^k]}^{U^k} f_u^k [p] du}$$

$$\bar{x}^k [p] = \int_{u=u^k[\lambda^k]}^{U^k} f_u^k [p] x^k [u; p] du$$

Finally, in order to find the Pareto-efficient allocation that can be realized in the largest number of ways, we can find the price system that has the highest entropy.

While individual agents are striving to maximize their individual utilities through the exchange process, from the maximum entropy point of view, the market is trying to maximize the disorder of the equilibrium allocation. Maximization of utility through voluntary exchange represents a constraint on the market process, not its objective. The Lagrangian shadow prices $\lambda [p]$ represent the entropy prices of the commodities, the amount of increase in entropy the market

could achieve at prices p if the total supplies of commodities were slightly increased.

There is an important relationship between the market prices p and the entropy prices $\lambda[p]$. The gradient of the entropy with respect to market prices can be directly calculated from the Lagrangian:

$$\frac{\partial H[p]}{\partial p} = -\lambda[p] \sum_{k=1}^r \left(\sum_{u=u^k[\lambda^k]}^{u^k} f_u^k \frac{\partial x^k[u;p]}{\partial p} \right)$$

Since with strictly increasing utility indicators the prices at any exchange equilibrium will be strictly positive, the vanishing of this gradient is a necessary (and since $H[.]$ is a concave function, a sufficient) condition for the maximization of entropy over market prices. But a well-known property of Hicksian demand functions is that the market prices p define the generically unique null space of the matrix $\frac{\partial x^k[u;p]}{\partial p}$. Since all the agents face the same market prices in an exchange equilibrium, the vanishing of the gradient of $H[.]$ at market prices p^* implies that $\lambda[p^*]$ is proportional to p^* .

Proposition 2: A necessary and sufficient condition for a price system p^* to maximize the entropy of the exchange equilibrium is that the entropy shadow prices $\lambda[p^*]$ be proportional to p^* .

This property may simplify the analysis of the comparative statics of maximum entropy exchange equilibria.

c. Visualization of the maximum entropy exchange equilibrium

The maximum entropy exchange equilibrium cannot be represented as the intersection of supply and demand curves (either Marshallian or Hicksian). Supply and demand curves represent agents' utility-maximizing bundles when they trade from their endowments at parametrically given prices. The

maximum entropy exchange equilibrium allows equilibrium prices to emerge from trading at disequilibrium prices, so that demand and supply curves cannot represent the corresponding trading process.

The Edgeworth box construction is also incapable of representing the maximum entropy exchange equilibrium, even when there are only two types of agents and two goods. The representation of an allocation as a point in the Edgeworth box rests on the implicit assumption that all the agents of each type receive the same consumption bundle, which is not true in a general exchange equilibrium.

At an exchange equilibrium all the agents of one type will have consumption bundles along the income-offer curve corresponding to the equilibrium prices. The agents will be spread out along the income-offer curve according to the distribution that arises from the maximum entropy programming problem.

4. Maximum entropy exchange equilibrium with a single, homothetic utility function

The example of a pure exchange economy where all the agents have the same homothetic utility function, but differ in their endowments will serve to make the abstract concept of the maximum entropy exchange equilibrium more concrete.

If the utility function of all the agents is homothetic, the gradient of the utility function at a consumption bundle (the marginal rate of substitution), which determines the terms at which the agent would exchange a small amount of any good for any other good, depends only on the proportions in which the agent consumes the goods, not on the scale of consumption. Since the arbitrated allocations are characterized by the requirement that all the agents have the same marginal rates of substitution, the assumption of a homothetic utility function assures that all the agents will be consuming goods in the same proportion in an arbitrated allocation, and hence, in order to achieve feasibility, in

the same proportions as the overall endowment of the economy, \bar{x} . There is only one possible exchange equilibrium price system, equal to the marginal rate of substitution vector, at the economy wide endowment, $p = \partial u[\bar{x}]$. Every agent's consumption in an exchange equilibrium is proportional to the total endowment.

Without loss of generality, choose a monotonic transformation of the utility function so that $u[0] = 0$, $u[\bar{x}] = 1$, and $u[\alpha x] = \alpha u[x]$ for all x . Then the utility of an agent at a feasible, arbitrated allocation is exactly equal to her share of the total endowment of the economy, since $u^i = u[\alpha^i \bar{x}] = \alpha^i u[\bar{x}] = \alpha^i$.

In order for an allocation to be individually preferred to the endowment for each type of agent, we must have $u^i \geq u[\alpha^{k[i]}]$ where $k[i]$ is the type of agent i . Thus the state space for this problem consists of a 1-dimensional assignment (u^i) , showing the utility of, or equivalently the proportion of the total resources consumed by, agent i .

We can characterize the feasible, arbitrated, individually improving allocations as vectors (u^1, \dots, u^n) with $u^i \geq u_{\min}^{k[i]}$, $\sum_{i=1}^n u^i = 1$.

Divide the utility line into small finite intervals of grain ϵ . The highest utility any agent can get in a feasible allocation is 1. We have $r+1$ constraints on the proportion of agents of each type occupying each utility segment, $\{f_u^k\}$:

$$\sum_{u=u[\epsilon^k]}^1 f_u^k = 1, \quad k = 1, \dots, r$$

$$\sum_{k=1}^r n_k \sum_{u=u[\epsilon^k]}^1 f_u^k = 1$$

The first constraints require that we place each agent in some individually improving utility state, and the second assures the feasibility of the allocation.

We can calculate the maximum entropy exchange equilibrium of this economy from the general first order conditions:

$$f_u^k = \frac{\exp[-\beta u]}{\sum_{u=u[\beta^k]} \exp[-\beta u]}$$

The average share (or utility) of an agent of type k is then:

$$\bar{u}^k = \sum_{u=u[\beta^k]} f_u^k u = \frac{\sum_{u=u[\beta^k]} \exp[-\beta u] u}{\sum_{u=u[\beta^k]} \exp[-\beta u]}$$

The single parameter β must be chosen to satisfy the feasibility constraint:

$$\sum_{k=1}^r w^k \frac{\sum_{u=u[\beta^k]} \exp[-\beta u] u}{\sum_{u=u[\beta^k]} \exp[-\beta u]} = \sum_{k=1}^r w^k \bar{u}^k = \frac{1}{n} \quad (C2')$$

It makes practically no difference whether the upper limits in the sums are made equal to 1 or to ∞ , since in order to meet the feasibility constraint the solution has to assign a tiny fraction of an agent to any u above 1.

Using the properties of sums of geometric series we have:

$$\sum_{u=u[\beta^k]} \exp[-\beta u] = \frac{\exp[-\beta u[\beta^k]]}{1 - \exp[-\beta]}$$

$$\sum_{u=u[\beta^k]} \exp[-\beta u] u = \frac{\exp[-\beta u[\beta^k]]}{(1 - \exp[-\beta])^2} (\exp[-\beta] + u[\beta^k] (1 - \exp[-\beta]))$$

Thus we have

$$\bar{u}^k = u[\bar{x}^k] + \frac{\exp[-\beta]}{1 - \exp[-\beta]}$$

From (C2'), then, we see that

$$\sum_{k=1}^r w^k \bar{u}^k = \sum_{k=1}^r w^k u[\bar{x}^k] + \frac{\exp[-\beta]}{1 - \exp[-\beta]} = \frac{1}{n}$$

If we define $\bar{u}_{\min} \equiv \sum_{k=1}^r w^k u[\bar{x}^k]$ as the average of the minimum utilities across different types, we see that:

$$\frac{\exp[-\beta]}{1 - \exp[-\beta]} = \frac{1}{n} - \bar{u}_{\min}$$

$$\bar{u}^k = \frac{1}{n} + (u_{\min}^k - \bar{u}_{\min})$$

Thus the average utility, or share, of type k is the average share for the whole economy, $\frac{1}{n}$, plus the deviation of type k 's minimum utility from the average of the minimum utilities for all the types.

The Walrasian competitive equilibrium for this economy would give each agent of type k a share equal to the ratio of the values of \bar{x}^k and \bar{x} at the equilibrium prices.

$$u_{\text{walrasian}}^{k[i]} = \frac{u_{\text{walrasian}}^{-k}}{u_{\text{walrasian}}^{-i}} = \frac{p_{\text{walrasian}} \bar{x}^k}{p_{\text{walrasian}} \bar{x}}$$

The two concepts predict quite different distributions of goods among agents of a given type, and different average consumption bundles for each type. The maximum entropy exchange equilibrium puts more weight on the agent's bargaining position when trading at disequilibrium prices, which is represented by her minimum utility, than on the value of her endowment at the equilibrium price system.

5. Walras' auctioneer and Maxwell's demon

The effort of individual agents to increase their utility through bilateral trading produces a large amount of information in that it greatly reduces the entropy of the equilibrium allocation. We could, for instance, think of

maximizing the entropy over all feasible allocations without the arbitrage and individual improvement constraints. The resulting allocation would have a higher entropy than the maximum entropy exchange equilibrium, reflecting the loss of order sustained by arbitrage through trade. Any constraint on the final allocation will tend to reduce the entropy of the allocation, and in that sense, any constraint represents the production of information.

The Walrasian competitive equilibrium constraint limits the equilibrium allocation in requiring that the values of consumption bundles be equal to the values of endowments at the equilibrium prices. The reduction in entropy corresponding to this constraint is very large. Since in Walrasian competitive equilibrium all agents of the same type have the same consumption bundle, the entropy of the Walrasian equilibrium is zero. The huge amount of information required to sustain this very unlikely allocation is reflected in the classic Walrasian theory by the assumption that all agents know the equilibrium prices before they enter into trade.

It has always been a puzzle for the Walrasian approach to explain just how this enormous amount of information might be generated. Walras' own explanation was to posit the existence of an "auctioneer", an agent whose only purpose was to generate equilibrium prices, outside of real trading time, without using any of the resources of the economy, and without any actual trading. There is a paradox inherent in this conception: Walras' auctioneer generates information without expending any resources.

In this respect Walras' auctioneer is like Maxwell's imaginary demon who controls an opening in an otherwise impermeable membrane separating two volumes of a fluid. As molecules of different energies approach the opening the demon cleverly allows only high energy molecules to pass into one compartment and low energy molecules to pass into the other compartment, thereby maintaining a temperature

difference between the two parts of the fluid which can be exploited in a heat engine to do useful work. Such a demon would make possible a perpetual motion machine, or equivalently, an endless source of costless useful energy. Walras' auctioneer performs an exactly analogous role in the exchange market, costlessly ordering an otherwise highly disordered system.

An alternative to Walras' auctioneer is Edgeworth's idea of allowing universal costless recontracting without trade among all subsets of agents, as in the analysis of the core of the economy (see Herbert Scarf and Gerard Debreu, 1962.) This mechanism has the same informational properties as Walras' auctioneer, since the core allocations of an exchange economy coincide in the limit of large numbers of traders with the Walrasian competitive equilibrium allocations.

There are many entropy-reducing devices in real markets, such as brokers, arbitrageurs, and advertising media, all of which function to reduce trading at disequilibrium prices. But in real markets, as in the real world of thermodynamics, these devices expend real resources, establishing a practical tradeoff between entropy reduction and real resources available for other uses.

Attempts to model disequilibrium trading (such as Frank Hahn and T. Negishi, 1962, Peter Diamond, 1984, and Peter Albin and the present author, 1991, to mention a random and far from exhaustive sample) have sought to specify the dynamics of the trading process, and to solve the resulting dynamic problem to predict the final allocation. The choice of a specific mechanism of trade inevitably appears arbitrary in these papers. The example of physics suggests that statistical considerations must be introduced in order to make progress on this type of dynamic problem.

6. Statistical welfare economics

Walrasian competitive equilibrium assures horizontal equality: agents who share the same tastes and endowments

receive the same consumption bundle (and utility) in the equilibrium allocation.

Exchange equilibria in general, including the maximum entropy exchange equilibrium, exhibit horizontal inequality. For example, in the exchange economy with common homothetic utility indicators, the agents of each type are spread out along the utility space, with the largest proportion receiving the minimal utility level. Differences in final consumption and welfare in Walrasian competitive equilibrium always correspond to differences in initial endowments, while agents with the same initial endowments can have quite different consumption and utility levels in exchange equilibria. Walrasian competition, because it rules out trading at disequilibrium prices, does not allow the market process to introduce differential gains for agents of the same type, while in general exchange equilibrium does allow for endogenous market contributions to inequality.

If real world markets generate an important horizontal inequalities endogenously, analyses of economic policy based on the model of Walrasian competitive equilibrium will ignore an important aspect of reality. Take, for instance, the issue of legislated price floors and ceilings. From a Walrasian point of view, these limitations on voluntary exchange are either irrelevant (if the constraints are not binding at the Walrasian competitive equilibrium) or inefficient (if they do bind, and prevent market prices from reaching their equilibrium). But this analysis depends on the prior assumption that all trades actually take place at equilibrium prices. If a significant amount of trading takes place at disequilibrium prices, price floors and ceilings can serve to protect agents against relatively disadvantageous trades, and thus to mitigate the endogenous horizontal inequality generated by the market.

7. Conclusion

The statistical approach to the analysis of phenomena involving a large number of essentially indistinguishable agents has been remarkably successful in physics, not least because it suggests immediate and direct connections between empirically observed system-wide averages and the underlying micro-properties of the system. It is difficult to exploit this methodological advantage in the context of the pure exchange model, however, because it is so difficult to interpret measured prices and quantities in ongoing productive economies as the outcome of once and for all establishment of equilibrium through trade. If these methods can be applied to more adequate dynamic models of production economies existing in time empirical-theoretical links should be easier to forge.

8. References

Albin, Peter, and Duncan K. Foley. 1991. Decentralized, dispersed exchange without an auctioneer: a simulation study. Journal of Economic Behavior and Organization, forthcoming.

Diamond, Peter A. 1984. A Search Equilibrium Approach to the Micro Foundations of Macroeconomics. Cambridge, MA: M.I.T. Press.

Hahn, Frank, and T. Negishi. 1962. A theorem on non-tatonnement stability. Econometrica 30 (July) 463-69.

Jaynes, E. T. 1978. Where do we stand on maximum entropy? Reprinted in R. D. Rosenkrantz, 1982, pp 210-314.

Jaynes, E. T. 1991. Probability theory--the logic of science. Unpublished ms.

Mirowski, Philip. 1990. More heat than light: Economics as social physics: Physics as nature's economics. Cambridge: Cambridge University Press.

Rosenkrantz, R. D. ed. 1982. E.T. Jaynes: Papers on Probability, Statistics and Statistical Physics. Dordrecht: D. Reidel.

Scarf, Herbert and Gerard Debreu. 1962. A limit theorem on the core of an economy. International Economic Review 4 (September), 235-46.