Decentralized Trade, Random Utility and the Evolution of Social Welfare*

KANDORI, Michihiro,† Roberto SERRANO‡ and Oscar VOLIJ§

Published in Journal of Economic Theory, 140(1), 328-338, 2008.

Abstract

We study decentralized trade processes in general exchange economies and house allocation problems with and without money. The processes are affected by persistent random shocks stemming from agents' maximization of random utility. By imposing structure on the utility noise term –logit distribution–, one is able to calculate exactly the stationary distribution of the perturbed Markov process for any level of noise. We show that the stationary distribution places the largest probability on the maximizers of weighted sums of the agents' (intrinsic) utilities, and this probability tends to 1 as noise vanishes.

JEL classification numbers: C79, D51, D71.

Keywords: decentralized trade, exchange economies, housing markets, long-run stochastic stability, logit model, social welfare functions.

*We thank an Associate Editor and three anonymous referees for useful comments, and Ryo Ogawa for his research assistance. Kandori is grateful for research support by JSPS (grant 17530138). Serrano thanks Fundación Banco Herrero, Universidad Carlos III, the National Science Foundation under grant SES-0133113, and Deutsche Bank for research support, and Universidad Carlos III and CEMFI in Madrid and the Institute for Advanced Study in Princeton for their hospitality. Volij thanks the Spanish Ministerio de Educación y Ciencia (project SEJ2006-05455) for research support.

†(Correspondent) Faculty of Economics, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan, e-mail: kandori@e.u-tokyo.ac.jp

‡Brown University, e-mail: roberto_serrano@brown.edu

§Iowa State University and Ben-Gurion University of the Negev, e-mail: oscar@volij.co.il

1 Introduction

This paper considers the allocation of indivisible durable goods through decentralized trading processes. A simple example is the allocation of N offices among N students. Even if the size of the problem, N, is relatively small, the number of possible allocations can be quite large. With ten students and offices, the number of allocations is about 3.6 million. We examine how successful decentralized trading processes, both with and without money, are in solving those complex combinatorial problems.

We consider a situation where agents randomly meet over time. When a group of agents meet, they exchange their goods in the following simple way. First, a new allocation for them is randomly proposed, and it is accepted if it provides a higher utility for all of them. Otherwise, the agents continue to hold their endowments.[1] When they assess the proposed allocation, we assume that their utility is affected by random shocks. The shocks can be interpreted as mistakes, or transitory changes in tastes attributed to noise; see, for example, the related notion of quantal response equilibrium in behavioral game theory (McKelvey and Palfrey [8]).

Incorporating random terms in utility functions has been found to be quite useful in econometric studies of discrete choice problems and behavioral game theory, and we employ one of the leading specifications, the *logit model*, for the distribution of the noise term. Thanks to the special structure of the model, we obtain the closed form solution of the stationary distribution, for *any level of noise*. In this respect, our work is built on the literature pioneered by Blume [3, 4], who identified a set of conditions which enables one to derive the closed form stationary distribution under logit noise. Our technical contribution is to show that a similar closed form can be obtained in a wider class of models, even when Blume's conditions are not satisfied.[2]

This approach contrasts with the traditional long-run stochastic stability methodology (see Kandori, Mailath and Rob [6] and Young [11]). The method identifies those states –allocations– in which the economy spends most of its time in the long run, when the noise in the system is made *negligible*. Negligible noise implies a fairly long waiting time to see the long run effects.

[1] Hence, agents in our model act myopically. Either bounded rationality or the arrival rate of trading opportunities being low, relative to the discount factor, might justify this.

[2] Section 4 provides a detailed discussion.

The present paper, in contrast, allows us to characterize the stationary distribution *for any level of noise*. Specifically, we show that, for any level of noise, the states that maximize a weighted sum of the agents' intrinsic utilities receive the largest probability in the stationary distribution.

Our result sheds light on the previous contribution by Ben-Shoham, Serrano and Volij [2]. They considered house allocation problems and found that, with vanishing noise, *the minimum envy allocation* is selected when serious mistakes are less likely. An agent's envy level is the number of other agents who have better houses, and the minimum envy allocation is the one that minimizes the aggregate envy level. We show that this somewhat mysterious result can be derived from a more general principle, namely, that evolutionary dynamics with logit noise maximize the aggregate utility level.[3]

Our results imply, in particular, that the most likely state is efficient. Note that, with no noise, our exchange processes may be stuck on an inefficient state. For example, when only *bilateral trades* are possible, the society may be stuck on an inefficient state where there is *no double coincidence of wants*.[4] In this respect, our decentralized trading processes resemble the algorithms to solve combinatorial optimization problems, where the process may get stuck at a local maximum. For the latter problems, random search algorithms, notably simulated annealing methods (see Aarts and Korst [1]) have been found quite effective. Just like randomness in simulated annealing helps to escape from a local maximum, so does the randomness in our trading process to ensure that the society is not stuck at an inefficient state.

The paper is organized as follows. Section 2 analyzes the dynamic model for discrete barter economies. Using the assumption of quasilinear utilities, Section 3 presents a version of the result that introduces monetary side payments that take place in a continuous money unit. The final section discusses related literature.

[3] See Section 4 again for details.

[4] Ben-Shoham, Serrano and Volij [2] showed that an *in*efficient state can be stochastically stable, when all mistakes are equally likely. Hence adding noise does not always help escape from an inefficient state. Our model provides a set of sufficient conditions for the noise term to knock out inefficient states.

2 Decentralized Barter: Exchange Economies

There are K durable and indivisible commodities in the economy. The set of agents is $N = \{1, \ldots, I\}$. Agent i's consumption set is $X_i \subset \{0, 1, 2, ...\}^K$. This allows for the possibility that an agent consumes an arbitrary number of units of each good, as in general exchange economies, or only one unit of one of the (heterogeneous) goods, as in house allocation problems. At time $t \in \{1, 2, \ldots,\}$ agent i holds a bundle of commodities denoted by $z_i(t)$. Although the individuals' holdings may change over time, the aggregate endowment of goods remains fixed, i.e. $\sum_{i \in N} z_i(t) = \overline{z}$. A *coalition* is a non-empty subset of agents. For any coalition $S \subset N$, a feasible allocation for S at time t is a distribution of their endowments at t. Thus, the set of feasible allocations for S at t is $A_S(z_S(t)) = \{z'_S \in \times_{i \in S} X_i \mid \sum_{i \in S} z'_i = \sum_{i \in S} z_i(t)\}$. In particular, the set of feasible allocations in the economy is given by $Z = A_N(\overline{z}) = \{z'_N \in \times_{i \in N} X_i \mid \sum_{i \in N} z'_i = \overline{z}\}$.

There is an exogenously given set of allowable coalitions, denoted $\mathbf{S} \subset 2^N$ that may meet and trade in each period. For example, when only pairwise meetings are possible (a particular case of our model), we have $\mathbf{S} = \{S \subset N \mid |S| = 2\}$. At period $t = 1, 2, \ldots$ a coalition $S \in \mathbf{S}$ is selected with probability $q(S) > 0$ (independent of time), and has the opportunity to reallocate their holdings of commodities. We assume that from any initial feasible allocation z, any feasible allocation z' can be reached through a series of feasible proposals by a finite sequence of allowable coalitions $S^1, \ldots, S^T \in \mathbf{S}$.

Suppose that, in the current period, a coalition $S \in \mathbf{S}$ is selected, and let $z_S \equiv z_S(t)$ be the allocation of goods for this coalition at the beginning of the current period. A new allocation for this coalition is chosen according to a probability distribution, which may depend on the current allocation, over the set of feasible allocations $A_S(z_S)$. We assume that there is certain symmetry in the proposal distribution.

Assumption 1 For any $z_S, z'_S \in A_S(\cdot)$, the probability that allocation z'_S is chosen when the current allocation is z_S is the same as the probability that allocation z_S is chosen when the current allocation is z'_S.

There are some instances where this requirement is naturally satisfied. For example, this assumption holds when proposals are completely random (a new allocation is drawn from the uniform distribution over the set of feasible allocations for the coalition). Another example is a house allocation problem with pairwise trade: Assumption 1 is satisfied if a pair of players, whenever they meet, always propose to exchange their houses.

We assume that agents' utilities are affected by random shocks, so that agent i's utility is given by

$$v_i(z_i) = u_i(z_i) + \eta_i(z_i), \tag{1}$$

where $u_i(z_i)$ and $\eta_i(z_i)$ stand for the *intrinsic utility* derived from the bundle z_i and *noise*, respectively. We assume that, when coalition S is formed, they adopt a *(myopic) unanimity rule*: when allocation z'_S is proposed instead of z_S, it is adopted if and only if $\forall i \in S$, $u_i(z'_i) + \eta_i(z'_i) \geq u_i(z_i) + \eta_i(z_i)$, with a strict inequality for at least one agent. We assume that the noise term has the following distribution.[5]

Assumption 2 The noise term is independently distributed over time and across agents according to the *type I extreme value distribution* (or *Gumbel distribution*) with *precision parameter* $\beta_i > 0$, whose cumulative distribution function F_i is given by

$$F_i(x) = \exp(-\exp(-\beta_i x - \gamma_i)), \tag{2}$$

where γ_i is a constant so that the resulting mean equals zero.

Note that agent i's preferences over z'_i and z_i depend on the random variable $\eta_i(z'_i) - \eta_i(z_i)$. The above assumption basically implies that it has a bell-shaped distribution which is *quite similar to a normal distribution*. One reason to consider this rather specific distribution is to obtain a tractable model, which approximates the normal noise model (this does not admit a closed form stationary distribution; see Remark 1 at the end of this section for more discussion). When the

[5] Since the noise terms can be negative, the marginal utility of goods can sometimes be negative. Consider, for example, two bundles of durable goods z and z', where z' is obtained by adding one bicycle to z. Even though the intrinsic marginal utility of a bike is positive ($u_i(z') - u_i(z) > 0$), the realized marginal utility may sometimes be negative ($[u_i(z') + \eta_i(z')] - [u_i(z) + \eta_i(z)] < 0$). This captures the situation where the agent *on average* would like a bike, but sometimes he is happy to give it up for various reasons (such as a transitory negative shock to his health conditions).

noise term $\eta_i(z_i)$ is distributed according to (2), it is known that the probability that agent i agrees to receive z_i' in exchange for z_i is given by

$$\Pr(v_i(z_i') > v_i(z_i)) = \frac{\exp[\beta_i u_i(z_i')]}{\exp[\beta_i u_i(z_i')] + \exp[\beta_i u_i(z_i)]}. \tag{3}$$

$>$From this formula it can be seen that, as $\beta_i \to \infty$, noise vanishes and the agent maximizes u_i without any error. This distributional assumption is what is behind the *logit model* in econometrics.

The above description defines a Markov process on the set of feasible allocations of the economy. At every period, the economy can transit from one allocation to another and, since we assumed that it is always possible to go from any allocation to any other through a finite sequence of feasible reallocations, the resulting Markov process is *irreducible*. Moreover, there is a chance that the state does not change, which makes the process *aperiodic*. For an irreducible and aperiodic process, there is a unique *stationary distribution* with the following two properties. Firstly, starting from any initial allocation, the probability distribution on period t allocations is known to approach that stationary distribution as $t \to \infty$. Secondly, the stationary distribution also represents the proportion of time spent on each state over an infinite time horizon. Our first result characterizes this stationary distribution.

Proposition 1 *In the barter model with random utility, the stationary distribution over the set of allocations is given by*

$$\mu(z) = \frac{\exp \sum_{i \in N} \beta_i u_i(z_i)}{\sum_{z' \in Z} \exp \sum_{i \in N} \beta_i u_i(z_i')}.$$

Before we present the proof, a few remarks are in order. First, the formula tells us that the stationary distribution is "exponentially proportional" to the *utilitarian social welfare function* $\sum_{i \in N} \beta_i u_i(z_i)$. In particular, the most likely states (for any level of noise) are the ones that maximize that social welfare. Second, recall that β_i is the precision parameter of agent i's noise term, meaning that a larger β_i implies a smaller level of noise. The formula is easiest to understand when we regard the noise term as the representation of mistakes; an agent who makes fewer mistakes (i.e., who has a higher β_i) has a higher weight in the long run distribution. Third, the stationary distribution is independent of the matching probabilities, represented by $q(s)$. Suppose that we have two agents with identical utility functions and precision parameters, and assume that one has more opportunities to trade than the other. Although one might expect that the one with more opportunities to trade does better than the other, in the long run they receive the same payoff distribution.

Proof. Let $\Pr(z, z')$ be the transition probability from z to z'. It is enough to show that

$$\mu(z) \Pr(z, z') = \mu(z') \Pr(z', z) \qquad \forall z, z' \in Z. \tag{4}$$

To see that this is sufficient, note that by summing both sides over all $z' \in Z$ we get

$$\mu(z) = \sum_{z' \in Z} \mu(z') \Pr(z', z) \qquad \forall z \in Z,$$

which means that μ is a stationary distribution. Equation (4) is what is known as the *detailed balance condition*, and it says that the probability inflows and outflows are balanced for any pair of states. Our symmetric proposal assumption 1 implies $\Pr(z, z') = 0 \Leftrightarrow \Pr(z', z) = 0$, so that (4) is satisfied in such a case. In the remaining case, the closed form formula of $\mu(z)$ implies that the detailed balance condition is satisfied if

$$\frac{\exp \sum_{i \in S'} \beta_i u_i(z_i)}{\exp \sum_{i \in S'} \beta_i u_i(z'_i)} = \frac{\Pr(z', z)}{\Pr(z, z')}, \tag{5}$$

where $S' \equiv \{i \in N \mid z'_i \neq z_i\}$ is the set of agents who have different bundles at z and z'. Now let us calculate the transition probabilities $\Pr(z, z')$ and $\Pr(z', z)$. Let $\mathbf{S}' \equiv \{S \in \mathbf{S} \mid S' \subset S\}$ be the set of feasible coalitions containing S'. Starting with z, the new allocation z' is obtained if and only if a coalition $S \in \mathbf{S}'$ is selected, proposal z'_S is made, and all members of S' prefer z'_i to z_i.[6] Recalling that $q(S)$ is the probability that coalition S is selected to make a proposal, and denoting by $r_{z_S}(z'_S)$ the probability that S proposes z'_S, we have, using (3),

$$\begin{aligned} \Pr(z, z') &= \sum_{S \in \mathbf{S}'} q(S) r_{z_S}(z'_S) \prod_{i \in S} \Pr(v_i(z'_i) > v_i(z_i)) \\ &= \sum_{S \in \mathbf{S}'} q(S) r_{z_S}(z'_S) \frac{\exp[\sum_{i \in S'} \beta_i u_i(z'_i)]}{H}, \end{aligned}$$

where $H = \prod_{i \in S'} \{\exp[\beta_i u_i(z'_i)] + \exp[\beta_i u_i(z_i)]\}$. Similarly, we have

$$\Pr(z', z) = \sum_{S \in \mathbf{S}'} q(S) r_{z'_S}(z_S) \frac{\exp[\sum_{i \in S'} \beta_i u_i(z_i)]}{H}.$$

[6] Agents in $S \setminus S'$ are proposed the same bundles as before, so they are indifferent between z'_S and z_S.

By our symmetric proposal assumption 1, we have $r_{z_S}(z'_S) = r_{z'_S}(z_S)$, and the condition (5) is satisfied. ■

Note that the detailed balance equation (4) fails when the proposal distribution does not satisfy assumption 1, as the proof shows: without this assumption, the clean closed form solution cannot be obtained.

Let us now examine how the stationary distribution changes with the level of noise. For simplicity, consider the symmetric case with $\beta_1 = \cdots = \beta_I = \beta$. As the level of noise decreases (i.e., as β increases), states with higher social welfare $\sum_{i \in N} u_i(z_i)$ receive higher probabilities. When noise is vanishing ($\beta \to \infty$), each term $\exp \beta \sum_{i \in N} u_i(z_i)$, $z \in Z$ diverges to infinity, but the one that corresponds to the maximizer of the social welfare $\sum_{i \in N} u_i(z_i)$ does so with the highest speed. Hence we have the following characterization.

Corollary 1 *In the barter model with random utility, if the noise is symmetric $\beta_1 = \cdots = \beta_I = \beta$, then as $\beta \to \infty$, the limiting stationary distribution places probability 1 on the set of allocations that maximize the sum of the agents' intrinsic utility functions.*

One can generalize the above corollary as follows: if for all $i \in N$ the noise parameter is $\beta_i = \lambda_i \beta$ for some $\lambda_i > 0$, then as $\beta \to \infty$, the limiting stationary distribution places probability 1 on the set of allocations that maximize the weighted utilitarian social welfare function $\sum_{i \in N} \lambda_i u_i(z_i)$.

Remark 1 In the discussion paper version (Kandori, Serrano and Volij [7]), we have conducted numerical simulations to compare our logit noise model with the one with normal noise. The sample paths in the two models are similar, although convergence to the efficient allocation appears to be slower in the normal case. This may come from the fact that the logit distribution has fatter tails, so that large shocks are more likely.

Remark 2 Assume now that the noise term enters the utility function in the following *multiplicative form*:

$$v_i(z_i) = u_i(z_i)\xi_i(z_i), \tag{6}$$

where $\log \xi_i(z_i)$ has type I extreme value distribution with parameter β_i. Then we can replace u_i with $\log u_i$ in the formula in Proposition 1 to obtain that the stationary distribution $\mu(z)$ is proportional to the weighted *Nash social welfare function* $\prod_{i \in N} u_i(z_i)^{\beta_i}$. The random utility functions of the original model are monotone transformation of the ones in this model (i.e., $\log u_i(z_i)\xi_i(z_i) = u_i'(z_i) + \eta_i(z_i)$, where $u_i'(z_i) = \log u_i(z_i)$ and $\eta_i(z_i) = \log \xi_i(z_i)$), so that they can merely be viewed as different representations of the same random preferences. When there is a clear-cut cardinal meaning attached to the intrinsic utility, however, we can view them as intrinsically different. For example, consider the case where $u_i(z_i)$ is interpreted as a von Neumann-Morgenstern (or Bernoulli) utility function, or the monetary value of (or willingness to pay for) z_i in the case where agent i has a quasi-linear utility function. The situation where willingness to pay is subject to additive noise is intrinsically different from the case where willingness to pay has multiplicative noise. We can then say that the most likely state in the long run maximizes the utilitarian social welfare function $\sum_{i \in N} \beta_i u_i(z_i)$ in the former case, while the Nash social welfare is maximized in the latter.

3 Trade with Money: House Allocation Problems with Side Payments

We now consider the case where indivisible goods are traded with monetary side payments. While the barter model of the previous section may be a good approximation of the office allocation in a department, in order to describe a housing market it would be more realistic to introduce monetary transfers. Specifically, we consider an economy with a set H of houses, and a set N of agents, each of whom occupies one house. Hence, the number of houses is the same as the number of agents: $|H| = |N|$. At each moment, agent i possesses one house and money, (z_i, m_i), where the real number m_i denotes agent i's money holdings. We take the standard partial equilibrium interpretation that m_i represents the flow activities not explicitly modeled.[7] Namely,

[7] An alternative formulation would be to treat money as one of the durable goods in the previous section's model. However, this formulation suffers from the problematic feature that agents derive flow utility from monetary balance, even though they do not spend it on goods and services. Our partial equilibrium formulation avoids such a problem.

it incorporates income flow and expenditure on other goods besides the house. In particular, we assume that there is enough income flow in each period so that agents afford the side payments associated with housing transactions. We also assume pairwise trade, where at each moment a pair of agents meet.

Each agent $i \in N$ is assumed to have quasi-linear utility:

$$\pi_i(z_i, m_i) \equiv v_i(z_i) + m_i = u_i(z_i) + \eta_i(z_i) + m_i. \tag{7}$$

As before, $\eta_i(z_i)$ is the random component of utility and it is distributed according to the type I extreme value distribution with a *common precision parameter* β for all $i \in N$. This turns out to be essential for the analysis in this section. Note that the quasi-linear utility implies no income effects, and this is one reason why we are able to "separate" the law of motion of durable goods from the dynamics of monetary transfers. This separation allows us, in turn, to apply finite Markov chain techniques to an economy with a perfectly divisible commodity (money).

In each period a pair of agents (i, j) is selected with probability $q(i, j) > 0$. Let $\underline{v}_i = u_i(z_i) + \eta_i(z_i)$ be agent i's willingness to pay for the house he currently owns, and let $\overline{v}_i = u_i(z_j) + \eta_i(z_j)$ be agent i's willingness to pay for agent j's house. The values $\underline{v}_j$ and $\overline{v}_j$ are defined in a similar way for agent j. Given the quasilinearity of preferences, there are potential gains from trade if $\overline{v}_i + \overline{v}_j > \underline{v}_i + \underline{v}_j$. Those gains are realized by a transfer p_i form j to i that satisfies $\overline{v}_i + p_i > \underline{v}_i$ and $\overline{v}_j - p_i > \underline{v}_j$. This is equivalent to say that (i) there is a non-empty *acceptable range of transfers* $R_{ij} \equiv (\underline{v}_i - \overline{v}_i, \overline{v}_j - \underline{v}_j)$ and (ii) $p_i \in R_{ij}$. In what follows, when $i < j$, we keep track of p_i, where a negative p_i represents a payment form i to j. We assume that the agents search over the possible range of transfers and detect the acceptable range R_{ij} with some probability $\mu_{ij}(R_{ij})$. When they detect such a range exists, they trade with probability one, and presumably they settle for a monetary transfer p_i in R_{ij}. We do not, however, need to assume any details of the bargaining procedure, or how the transfer is actually determined.

Assumption 3 There is a probability measure μ_{ij} on $(-\infty, \infty)$, such that agents i and j $(i < j)$ exchange their houses with probability $\mu_{ij}(R_{ij})$. Furthermore, μ_{ij} is symmetric: for all intervals $[a, b]$, $\mu_{ij}([a, b]) = \mu_{ij}([-b, -a])$.

Note that the probability measure μ_{ij} does not necessarily have full support. One may assume that the support of μ_{ij} is an appropriately bounded interval so that the monetary transfers are compatible with the implicitly assumed income flow. Note also that we assumed that bargaining is *in*efficient in the sense that agents may sometimes fail to trade even if R_{ij} is non-empty. In a remark at the end of this section, we will discuss what happens if one assumes efficient bargaining.

A *house allocation* is an assignment $(z_i)_{i\in N}$ of the houses in H to the agents in N. Note that if we ignore money holdings, the above trading process defines an irreducible Markov chain on the set of house allocations. We are interested in the invariant distribution induced by our trading process on the set of house allocations. Since after each pairwise meeting the probability of trade depends only on the measure of the acceptable range of transfers, and not on the actual transfer negotiated by the traders, the invariant distribution does not depend on the particular bargaining procedure the traders use. The following proposition shows that the invariant distribution is also independent of the measures μ_{ij}.

Proposition 2 *Under assumption 3, the stationary distribution for the allocation of houses is given by*

$$\mu^z(z) = \frac{\exp\left[\beta \sum_{i\in N} u_i(z_i)\right]}{\sum_{z'\in Z} \exp\left[\beta \sum_{i\in N} u_i(z'_i)\right]}.$$

The proof can be found in the appendix. The following remarks are in order:

Remark 3 The result can be extended to an exchange economy in which there are K indivisible goods (apart from money) and where an agent can hold any subset of the indivisible goods. To do this, as in Section 2, one needs to assume that the proposal distribution in each meeting is "symmetric." The result can also be extended to a process in which coalitions, not only pairs, trade, where side payments in coalition S, $p = (p_i)_{i\in S}$ satisfies $\sum_{i\in S} p_i = 0$.

Remark 4 If one assumes that the bargaining outcome is always fully efficient (with respect to the noise-perturbed utility), the closed form solution in Proposition 2 cannot be obtained. Under this assumption, agents i and j exchange their houses if doing so increases the total surplus;

$$[u_i(z_j) + \eta_i(z_j)] + [u_j(z_i) + \eta_j(z_i)] > [u_i(z_i) + \eta_i(z_i)] + [u_j(z_j) + \eta_j(z_j)].$$

If the sum of the noise terms $\eta_i(\cdot) + \eta_j(\cdot)$ had extreme value distribution, a simple modification of our argument in the previous section would yield the same stationary distribution as in Proposition 2 (and this would correspond to the one-player version of Blume's [4] model). However, this is not the case, since $\eta_i(\cdot) + \eta_j(\cdot)$ does *not* have extreme value distribution, even though $\eta_i(\cdot)$ and $\eta_j(\cdot)$ do.

4 Related Work

Our work generalizes a result due to Ben-Shoham, Serrano and Volij [2]. In that paper, only pairwise trade in the house allocation problem without money is considered, and the persistent shocks are "mistakes" in decision-making. In particular, they assume that, when an agent has his kth best house, the probability of accepting her mth best house ($m > k$) has the order of ε^{m-k}, where $\varepsilon \in (0, 1)$ is a small number. This is a particular formulation of mistake probabilities, where more serious mistakes are less likely. They showed that, when the randomness is vanishingly small (as $\varepsilon \to 0$), the allocation that *minimizes envy* is selected in the long run. Agent i's *envy level* is the number of people who have better houses than agent i (according to i's preferences). The *envy* in the society is the sum of individual agents' envy levels. The current paper shows that there is a more general mechanism at work operating behind the Ben-Shoham *et al* result. First, we note that their specification of noise can be related to the logit model. Let N be the number of houses/agents and let us assume $u_i(z_i) = 1, 2, \ldots, N$, where N is the utility of the best house. A straightforward calculation shows that one can obtain their mistake probabilities, when we add the logit noise term to this utility function. Second, one can see that the envy is equal to $\sum_i (N - u(x_i))$, and minimizing this expression is equivalent to maximizing the utilitarian social welfare $\sum_i u(x_i)$. We have found that the driving force of their result is that the logit noise model maximizes the utilitarian social welfare (and this is true for any specifications of utility functions).

The dynamic adjustment processes with logit noise, leading to a specific closed form stationary distribution ("Gibbs distribution"), has been studied extensively in the statistical mechanics and simulated annealing literature (Aarts and Korst [1]). Blume [3, 4] pioneered in applying this technique to game theory;[8] see also Durlauf [5] for broader applications in economics. Blume obtained a closed form expression of the stationary distribution when the following three conditions are satisfied: (i) players play a *potential game* (i.e., each player's best reply function is the same as in a hypothetical game in which players have an identical payoff ("potential")), (ii) at each moment in time, only *one player* can adjust and (iii) the random noise terms in players' utility functions have logit (or extreme value) distribution with *identical precision parameter* β. Blume [4] showed, under those assumptions, that the stationary distribution is the *Gibbs distribution*

$$\frac{\exp \beta P(a)}{\sum_{a' \in A} \exp \beta P(a')}, \tag{8}$$

where P is the potential, and A is the set of strategy profiles. A number of papers have followed the approach, including these: Young and Burke [12] presents an application of Blume's result to the geographical distribution of agricultural contracts in Illinois, Sandholm [10] applies it to Pigouvian pricing under externalities, and Myatt and Wallace [9] to a class of collective decision problems that includes the private provision of a public good. All those papers are special cases of Blume's work, where the aforementioned three conditions hold. In contrast to these existing works, the present paper shows that a similar technique, leading to a similar closed form expression, can be applied to a wider class of situations.

An essential, common feature of Blume's work, our paper, and the simulated annealing literature is that the stationary distribution is determined by the detailed balance condition (4). All three employ logit noise and the symmetry of possible changes of strategies (see our Assumption 1). They differ, however, in the additional assumptions that are employed to derive the detailed balance condition. Blume employed the aforementioned three extra assumptions (i)-(iii), but our work shows that the detailed balance condition can be satisfied even when none of those holds. (See Remark 4 in Section 3, which explains one aspect of the difference in greater detail.)

Appendix

Proof of Proposition 2. For each pair of traders i and j ($i < j$), let μ_{ij} be the probability measure specified in Assumption 3. Consider the following rather specific bargaining procedure that agents i and j may use to determine whether they trade houses and, if they do so, the transfer

[8] Blume [3] considers local interaction models, while Blume [4] contains a result for general K-player games.

that j will make to i. A price $p_i \in (-\infty, \infty)$ is randomly chosen according to the probability measure μ_{ij}. If the realized price falls inside the acceptable range, that is if $p_i \in R_{ij}$, then they trade houses and i receives p_i from j. If the price falls outside the acceptable range, there is no trade and no transfers are made. It is clear that if traders follow this procedure, the probability of trade is given by $\mu_{ij}(R_{ij})$. Therefore, Assumption 3 is satisfied. Since the law of motion of house allocations only hinges on Assumption 3, and it is independent of how transfers are actually made, if we figure out the stationary distribution corresponding *to this particular bargaining procedure*, we will have figured it out for all bargaining procedures that induce a probability of trade given by the measure μ_{ij}. This is what we are going to show. According to our bargaining procedure, when p_i is proposed, agents i and j agree to trade if and only if

$$u_i(z_j) + \eta_i(z_j) + p_i > u_i(z_i) + \eta_i(z_i), \text{ and} \tag{9}$$

$$u_j(z_i) + \eta_j(z_i) - p_i > u_j(z_j) + \eta_j(z_j). \tag{10}$$

Then, as the random utility shocks $\eta_i(\cdot)$ and $\eta_j(\cdot)$ have extreme value distributions, condition (9) is satisfied with probability

$$\frac{\exp(\beta(u_i(z_j) + p_i))}{\exp(\beta(u_i(z_j) + p_i)) + \exp(\beta u_i(z_i))}. \tag{11}$$

Similarly, condition (10) is satisfied with probability

$$\frac{\exp(\beta(u_j(z_i) - p_i))}{\exp(\beta(u_j(z_i) - p_i)) + \exp(\beta u_j(z_j))}. \tag{12}$$

Hence, given p_i, trade occurs with the product of the above probabilities, which is equal to

$$\frac{\exp[\beta(u_i(z_j) + u_j(z_i))]}{H(p_i)}. \tag{13}$$

Here, $H(p_i)$ is the product of the denominators of (11) and (12). Note that the common precision parameter β of the logit noise terms in utility is essential to eliminate p_i from the numerator. For the argument below, we also calculate the probability of the reverse inequalities of (9) and (10), i.e., $(1 - (11)) \times (1 - (12))$, which is equal to

$$\frac{\exp[\beta(u_i(z_i) + u_j(z_j))]}{H(p_i)}. \tag{14}$$

Suppose that now agents i and j possess z_j and z_i respectively, and they meet and come up with price $-p_i$. In that event, trade occurs if

$$u_i(z_j) + \eta_i(z_j) < u_i(z_i) + \eta_i(z_i) - p_i, \text{ and} \tag{15}$$

$$u_j(z_i) + \eta_j(z_i) < u_j(z_j) + \eta_j(z_j) + p_i. \tag{16}$$

Note that those are the reverse inequalities of (9) and (10), which hold with probability (14). The rest of the proof is similar to that of Proposition 1, so we only provide a sketch. Analogously to condition (5) in the proof of Proposition 1, we only need to show

$$\frac{\exp[\beta(u_i(z_i) + u_j(z_j)]}{\exp[\beta(u_i(z_j) + u_j(z_i)]} = \frac{\Pr(z', z)}{\Pr(z, z')}. \tag{17}$$

Using (13) and (14), the transition probabilities are given by

$$\Pr(z, z') = \int q(ij) \frac{\exp[\beta(u_i(z_j) + u_j(z_i))]}{H(p_i)} d\mu_{ij}(p_i) \text{ and}$$

$$\Pr(z', z) = \int q(ij) \frac{\exp[\beta(u_i(z_i) + u_j(z_j))]}{H(p_i)} d\mu_{ij}(-p_i).$$

By our symmetry assumption $\mu_{ij}(p_i) = d\mu_{ij}(-p_i)$, the crucial condition (17) is satisfied, which completes the proof.

References

[1] E. Aarts and J. Korst, *Simulated Annealing and Boltzmann Machines*, Chichester: John Wiley & Sons, 1989.

[2] A. Ben-Shoham, R. Serrano and O. Volij, The evolution of exchange, *J. Econ. Theory* **114**, (2004), 310-328.

[3] L. Blume, Statistical mechanics of strategic interaction, *Games Econ. Behav.* **5**, (1993), 387-424.

[4] L. Blume, "Population games," in B. Arthur, S. Durlauf, and D. Lane (eds.), *The Economic System as an Evolving Complex System II*, Reading, MA:Addison-Wesley, 1997.

[5] S. N. Durlauf, "Statistical mechanics approaches to socioeconomic behavior", in B. Arthur, S. Durlauf, and D. Lane (eds.), *The Economic System as an Evolving Complex System II*, Reading, MA:Addison-Wesley, 1997.

[6] Kandori, M., G. Mailath, and R. Rob, Learning, mutations and long run equilibria in games, *Econometrica* **61**, (1993), 29-56.

[7] Kandori, M., R. Serrano and O. Volij, "Decentralized trade, random utility and the evolution of social welfare," Working Paper 2004-06, Department of Economics, Brown University, 2004.

[8] R. D. McKelvey and T. R. Palfrey, Quantal response equilibria for normal form games, *Games Econ. Behav.* **10**, (1995), 6-38.

[9] D. Myatt and C. Wallace, "An evolutionary justification for thresholds in collective-action problems," Mimeo, Department of Economics, University of Oxford, 2006.

[10] W. H. Sandholm, Pigouvian pricing and stochastic evolutionary implementation, *J. Econ. Theory*, in press.

[11] H. P. Young, The evolution of conventions, *Econometrica* **61**, (1993), 57-84.

[12] H. P. Young and M. Burke, Competition and custom in economic contracts: a case study of Illinois agriculture, *Amer. Econ. Rev.* **91**, (2001), 559-573.