Equilibrium Concepts for Social Interaction Models

Larry Blume and Steven Durlauf[1]

1. Introduction

This paper fragment describes the beginning of an attempt to understand the relationship between several different binary social interaction models found in the literature. These include the mean field models of Brock and Durlauf (1997), conventional Nash equilibrium, that which we call a Gibbs equilibrium, of which Glaeser, Sacerdote and Scheinkman (1997) is an example, and stationary distributions of the dynamic strategy revision models of Blume (1993), Young (1993) and Kandori, Mailath and Robb (1993). One principle finding is that while the mean-field, Nash and dynamic equilibria always exist, Gibbs equilibria typically do not. Their existence depends on a great deal of agent homogeneity. Another finding is that, as the interaction radius grows, the mean-field and Gibbs equilibria (when they exist) converge. Also, when Gibbs equilibria exist, they are dynamic equilibria. The relationship between dynamic equilibria and mean-field equilibria when Gibbs equilibria do not exist is unknown as of this writing. However it is known that the short-run dynamics of strategy revision models with global interaction is well approximated over finite time horizons by the solution path to a differential equation suggested by the Brock-Durlauf mean field model when the population is large. This suggests that the invariant distribution for the dynamic model converges to mass on one or more of the mean-field equilibria as the population size grows.

2. The structure of interactions-based models

2.1. interdependent decisionmaking

The underlying logic of interaction models is straightforward. The object of the exercise is to understand the behavior of a population of economic actors rather than that of a single actor. The focus of the analysis is the externalities across actors. These externalities are the source of the social interactions. They are taken to be direct. The decision problem of any one actor takes as parameters the decisions of other actors. Hence the interactions approach treats aggregate social behavior as a statistical regularity of the individual interactions. A second feature of these models is that individual behavior is not as tightly modeled as it is in traditional economic equilibrium models. Individual choice is guided by payoffs, but has a random component. This randomness can be attributed to individual-specific variables not observed by the modeler, or to some form of bounded rationality.

Random choice has been significant both theoretically and empirically, and externalities are certainly not new, but the combination of the two along with an focus analytical

---

[1] Cornell University and the University of Wisconsin, respectively

focus on population behavior gives rise to new and interesting economic phenomena. These systems are highly non-linear and have multiple steady states. The system response to shocks can be quite complex.

Interaction models typically specify explicitly a probability density characterizing individual behavior conditional on exogenous characteristics which can be both global and individual-specific, and an interaction structure, a specification of who affects whom. To fix ideas, consider first a population of actors facing identical choices in which no externalities are present. Conditional on the exogenous characteristics X_i, the contemporaneous behavior of each actor is independent of the behavior of the rest of the population. Thus the joint conditional probability distribution characterizing population choices can be factored so that

$$\mathrm{Prob}(\omega_1, \ldots, \omega_I | X_1, \ldots, X_I) = \prod_{i=1}^{I} \mathrm{Prob}(\omega_i | X_i) \tag{1}$$

where ω_i is actor i's choice, X_i is her vector of individual characteristics. Conditional independence allows one to characterize aggregate behavior of the population through laws of large numbers. For example, when the probability distributions on the right hand side are identical, average behavior of the population is described by mean individual behavior. That is,

$$\frac{1}{I}\sum_{i=1}^{I} \omega_i \approx \frac{1}{I}\sum_{i-1}^{I} E(\omega_i | X_i)$$

for large I.

When individual decisions are contemporaneously interdependent, however, the factorization in equation (1) may fail to exist. The presence of the cross-actor externality introduces a strategic element into decisionmaking which needs to be incorporated into the specification of the model. One modelling strategy which generalizes equation (1) in a straightforward fashion is to explicitly model the externality in the factorization:

$$\mathrm{Prob}(\omega_1, \ldots, \omega_I | X_1, \ldots, X_I, Z_1, \ldots, Z_I) = \prod_{i=1}^{I} \mathrm{Prob}(\omega_i | X_i; Z_i), \tag{2}$$

$$Z_i = \mathrm{Prob}(\omega_{-i} | \mathcal{I}_i) \tag{3}$$

where $\mathcal{I}_i$ is player i's information. One example of this formulation is Nash equilibrium. Suppose that there is no variation in the variables X_i or, alternatively, that $\mathcal{I}_i = (X_1, \ldots, X_I)$. The strategies available to player i are $\omega_i \in \Omega_i$, and

$$\mathrm{Prob}(\omega_i | Z_i) = 0 \quad \text{if } \omega_i \in \operatorname{argmax} E(u_i(\omega_i, \omega_{-i}) | \mathcal{I}_i).$$

This extends to games of incomplete information where X_i is player i's private signal. But this general framework also encompasses other choice models, some of which we now describe.

2.2. A general model

To see how interactions-based methods work we will discuss some examples based on ideas in Brock (1993), Blume (1993, 1997), Brock and Durlauf (1997) and Durlauf (1997), which study binary choice decisions in the presence of social interactions. These examples illustrate hos a choiced-based model with standard economic assumptions can naturally produce an interactive environment. The general framework introduced below has been used to study out-of-wedlock births and high school dropout rates, decisions in which individuals are more likely to conform to the behavior of their "reference group".[2]

Formally, consider a population of I individuals. Suppose that each individual chooses one of two actions, labeled -1 and $+1$. Suppose that each individual's utility is quadratic in her action and the action of others. Each individual has knowledge of the mean of the average action. That is, Z is given by equation (3). Then conditional upon her information, actor i's expected utility can be written (after renormalization) as

$$V(\omega_i, X_i, Z_i) = h(X_i,)\omega_i - E\big(\sum_j J_{i,j}(X_i, X_j)(\omega_i - \omega_j)^2 | X_i; Z_i\big) + \epsilon(\omega_i)$$

This specification can be decomposed into a private component, $h(X_i)\omega_i + \epsilon(\omega_i)$, and the interaction effect, $E\big(\sum_j J_{i,j}(X_i, X_j)(\omega_i - \omega_j)^2 | X_i; Z_i\big)$. The private component can be further decomposed (without loss of generality) into its mean, $h(X_i)\omega_i$, and the stochastic deviation $\epsilon(\omega)$. The terms $J_{i,j}(X_i, X_j)$ is a a measure of the disutility of non-conformance. When the $J_{i,j}$ are all positive there is an incentive to conform. The presence of positive conformity effects gives rise to multiple equilibria and interesting dynamics. The random terms are independent, and assumed to be distributed according to the extreme value distribution with parameter $\beta(X_i)$. That is,

$$\text{Prob}\big(\epsilon(-1) - \epsilon(1) < z\big) = \frac{1}{1 + \exp(-\beta(X_i)z)}, \qquad \beta(\,\cdot\,) \geq 0$$

This model reduces to a particular instance of the standard binary choice framework when there are no interaction effects, when $J_{i,j} \equiv 0$.

From this distribution the individual choice probabilities can be computed.

$$\text{Prob}(\omega_i = 1 | X_i, Z) \;=\; \frac{1}{1 + \exp -2\beta\Big(h(X_i) + \sum_j E\big(J_{ij}(X_i, X_j)\omega_j | X_i; Z_i\big)\Big)} \tag{4}$$

---

[2] REFERENCES?

The conditionals can be multiplied to construct the joint distribution of play conditional upon Z and the vector X of actors' characteristics. The distribution is of the form

$$\mathrm{Prob}(\omega) \sim \prod_i \exp\bigl(h(X_i) + \beta(X_i) \sum_j E\bigl(J_{i,j}(X_i, X_j)\omega_i\omega_j | X_i; Z_i\bigr) \tag{5}$$

Finally, equilibrium requires equation (3), that actors know the correct mean population play.

2.3. The Game-Theoretic Limit

As the parameter β becomes large, the probability distribution on player i's choices puts converges to point mass on the set of best responses. In some cases the limit joint action distribution can be interpreted as a Nash equilibrium. Fix a vector X of characteristics, and suppose that $I_i = X$. Suppose that for this X, $J_{ij} = J_{ji}$ and $h_i = h_j = h$, and that there is a number J such that $J_{ij} = J$ or $J_{ij} = 0$. The strategic situation is as follows. if $J_{ij} \neq 0$, then players i and j are involved in a symmetric strategic interaction. The return to player i from choosing ω_i is the sum of the expected payoffs across all strategic interactions in which i is a partner. That is, $V(\omega_i, X_i, Z_i) - \epsilon(\omega_i)$ is the expected utility to playing ω_i when the other members of the population are playing according to the distributions Z_j. Note that the payoffs to every two-person game can be written as we do by identifying one strategy with $+1$ one with -1, and then writing the payoffs as in the non-random component of V. When the $\beta = \infty$ limit has a solution that is concentrated on a single strategy profile, that strategy profile is a pure strategy equilibrium of the population game we have just described. In fact, for generic parameter values (h and J), if the population game has a symmetric pure strategy Nash equilibrium, then for β large enough there is a solution μ_β to equation (5) such that the μ_β converge to point mass on that Nash equilibrium strategy profile.

That being said, we are more interested in the specifications $\beta < \infty$, and in situations of stratategic interaction that do not necessarily set up neatly as population games. Fortunately this model can be solved directly for a moderately rich class of specifications of the the private valuations, conformity and error distribution.

2.4. The general model with limited information

In the remainder of this section and in those to follow, we will assume that $\mathcal{I}_i = X$. Thus uncertainty due to incomplete information about the characteristics of others will be ignored. Consequently the parameters X_i will be suppressed in the presentation.

2.4.1 Uniform global interactions

One solvable class of models is studied in Brock and Durlauf (1997). Suppose that all actors are homogeneous. There are no individual characteristics. Suppose too that each actor interacts with every other actor, so that

$$J_{i,j} = \begin{cases} J/(I-1) & \text{if } i \neq j; \\ 0 & \text{if } i = j. \end{cases}$$

The I in this expression makes the marginal rate of substitution between the private and social components of preference independent of population size.

It is clear in this specification that, to the extent actors wish to conform, they attempt to match the average of choice in the population. If the population mean were m, equation (5) would be

$$\mathrm{Prob}_{I,\beta,h,J}(\omega; m) \sim \exp\{-\beta H_{I,h,J}(\omega; m)\}$$

where

$$H_{I,h,J}(\omega; m) = \sum_i h\omega_i + \beta J \omega_i m$$

Because of homogeneity of the actors, the equilibrium condition (3) requires that

$$m = E(\omega_i)$$

where the expectation is taken with respect to the distribution $\mathrm{Prob}_{I,\beta,h,J}(\omega; m)$. A computation from equation (4) shows that m is any solution to

$$m = \tanh \beta(h + Jm) \tag{6}$$

Equation (6) is well-known in the world of statistical physics. This special case has an important physical interpretation, and is known as the Curie-Weiss model of magnetization. The following characterization of the solutions to (6) is well-known:

Theorem: Multiple versus unique solutions.

i. If $\beta J \leq 1$ and $h = 0$, $m = 0$ is the unique solution to (6).

ii. If $\beta J > 1$ and $h = 0$ there are three solutions: $m = 0$ and $m = \pm m^*(\beta J)$. Furthermore, $\lim_{\beta J \to \infty} m^*(\beta J) = 1$.

iii. If $h \neq 0$ and $J > 0$ there is a threshold $C(h) > 0$ ($= +\infty$ if $h \geq J$) such that

- *a.* for $\beta h < C(h)$, there is a unique solution, which agrees with h in sign; and
- *b.* for $\beta h > C(h)$ there are three solutions, only one of which agrees with h in sign. Furthermore, as β becomes large the extreme solutions converge to ± 1.

iv. If $J < 0$ there is a unique solution which agrees with h in sign.

This theorem illustrates both the nonlinearities and the multiple steady states which are the hallmarks of interacting systems. The model is nonlinear with respect to a change in h, the private component of preference, on the mean behavior m of the population. Indeed, the effect of a change in h may be to increase the number of equilibria, which will exceed one when the strength of interactions is great enough.

The parameter m is of interest to the modeler as well as to the actors. Because this model preserves the factorization of the joint distribution of choices into the product of the distribution of individual choices, a strong law guarantees that m is approximately the (sample) average choice when I is large.

The underlying strategic situation corresponds to a coordination game played by a population of opponents. The strategy choice $+1$ (-1) is risk-dominant if $h \geq 0$ $(h \leq 0)$. As β becones large, the two extreme solutions converge to the pure strategy Nash equilibria.

2.4.2 Uniform local interaction

An alternative to global interaction models are models with a spatial structure. Actors have an address, and care about the behavior of only their neighbors. Schelling (1969, 1983) explored racial clustering in neighborhoods in a model of this type. He analyzed the consequences of families having even a slight preference for immediate neighbors like themselves. The consequences of individuals caring only about the identity of their immediate neighbors is a global pattern of racial segregation in housing. Local interaction models have achieved some popularity in evolutionary game theory in the work of Blume (1993), Ellison (1993), Binmore Samuelson Vaughn (199?) among others. Glaeser, Sacerdote and Scheinkman (1996) have taken local interaction models to data, and Glaeser and Scheinkman are reporting on that work here. Local interaction can be introduced to the Brock-Durlauf model by locating the actors on a graph and letting $J_{i,j} = 0$ if actors i and j are not neighbors.

Suppose again that all I actors have identical values of β and h. They are the vertices of a connected graph. Let n_i denote the number of edges with one endpoint at i; that is, n_i is the number of neighbors of i. Let $J_{i,j} = J/n_i$ if j is a neighbor of i, and 0 otherwise. Then from (4)

$$m_i = \tanh \beta \biggl(h + \frac{J}{n_i} \sum_j m_j) \biggr) \qquad i = 1, \ldots, I \tag{7}$$

In a symmetric equilibrium, $m_i = m_j$ and equations (7) reduce to equation (6), so the Theorem on multiple versus unique solutions holds for this model if the word solution is replaced with the phrase symmetric solution. However, there may also be asymmetric equilibria. For example, suppose there are five actors on a line and $J = 1$. For $\beta > \sqrt{2}$ there are equilibria where $m_1 = -m_5$, $m_2 = -m_4$, $m_3 = 0$ and the means are negative to the left of actor 3. To see this, suppose that $m_3 = 0$. Then $m_2 = \tanh \beta m_1/2$ and $m_1 = \tanh \beta m_2$. Consequently, m_1 must solve

$$m_1 = \tanh\bigl(\beta \tanh(\beta m_1/2)\bigr).$$

The right hand side maps $[-1, 1]$ into itself, is strictly increasing, and $m_1 = 0$ is a solution for all β. Furthermore, the derivative of the right hand side with respect to m_1 at $m_1 = 0$ is $\beta^2/2$. For $\beta > \sqrt{2}$ this slope exceeds 1, and so there must be two more solutions. Symmetry around 0 implies one is the negative of the other. Set m_1 equal to the negative solution. Now solve for m_2. Since m_2 is the hyperbolic tangent of a negative number, it is negative. Let $m_5 = -m_1$. We can see from the symmetry of the equations that $m_4 = -m_2$. Finally, $m_3 = \tanh\big(\beta(m_2 + m_4)/2\big) = \tanh 0$ which equals 0. For instance, when $\beta = 2\sqrt{2}$, $m_1 = -0.987$ and $m_2 = -0.884$. Some graphs allow only symmetric solutions, such as a circle and the completely connected graph.

2.4.3 Actor heterogeneity — private valuations

Our examples so far all assume actors are homogeneous. The next example reintroduces X_i. Suppose that each actor's private effect on utility is characteristic dependent. Specifically, suppose that the h_i are mean-0 i.i.d. random variables taking on the values ± 1. Assume global interaction as in the first example. From equation (4) we can compute that the conditional mean of ω_i is[3]

$$E(\omega_i | h_i) = \tanh \beta(h_i + Jm)$$

where m is the unconditional mean of each actor. Consequently, in equilibrium

$$\begin{aligned} m_i &= \tanh \beta(h_i + Jm) \qquad \text{for } i = 1, 2. \\ m &= m_1 + m_2 \,. \end{aligned}$$

To solve this model, aggregate the two m_i-equations to get

$$m = \frac{1}{2} \tanh \beta(1 + Jm) + \frac{1}{2} \tanh \beta(-1 + Jm)\,. \tag{8}$$

Any equilibrium solves (8). Conversely, if m solves (8), let $m_i = \tanh \beta(h_i + Jm)$. Both of these numbers are between -1 and 1, and by definition $m = (m_1 + m_2)/2$, so m_1 and m_2 solve the model. Consequently, to find equilibria it suffices to solve just (8).

Again the right had side maps $[-1, 1]$ into itself, is increasing, and symmetric. Consequently $m = 0$ is a solution. The following can be shown: There are two critical values of J, $J_1 < J_2$. For all $J < J_1$ the solution is unique. For all $J_1 < J < J_2$ there is a constant β_J such that for $\beta < \beta_J$ there is a unique solution while for $\beta > \beta_J$ there are five solutions. For $J_2 < J$ there are constants $0 < b_J < B_J$ such that for $\beta < b_J$ the solution is unique, for $b_J < \beta < B_J$ there are three solutions, and for $B_J < \beta$ there are five solutions. As β diverges, the non-zero solutions converge to ± 1 when there are three solutions, and when there are five solutions the two additional solutions converge to $\pm 1/J$.

When additional values of h are possible, the picture gets more complicated. If h can take on n possible values, then there can be anywhere between 1 and $2n + 1$ equilibria.

---

[3] Here we study only symmetric equilibria, where all actors have the same map from h_i to $!_i$. We have not investigated the possibility of non-symmetric equilibria.

2.4.4 Actor heterogeneity — social valuations

When the conformity effect varies across the population the model can look quite different. Suppose that the interaction coefficient J can take on two values: $J_1 < 0 < J_2$. Furthermore, assume $h > 0$. then for each actor,

$$E(\omega_i|J_i) = \tanh\beta(h + J_i m)$$

where m is the population mean. Suppose fraction p_i of the population has interaction coefficient J_i for $i = 1, 2$. Then equilibrium occurs where

$$\begin{aligned} m_i &= \tanh\beta(h + J_i m) \qquad \text{for } i = 1, 2. \\ m &= p_1 m_1 + p_2 m_2 \,. \end{aligned}$$

Just as in the previous example, it suffices to solve the aggregate equation

$$m = p_1 \tanh\beta(h + J_1 m) + p_2 \tanh\beta(h + J_2 m)\,. \tag{9}$$

One of several different cases occurs, depending upon the relative values of the J_i and h. The most interesting cases occur when for both i, $|J|_i > h$. The example is best understood by examining the right hand side when β is large. The interval $[-1, 1]$ divides into three parts: $[-1, -h/J_2]$, $[-h/J_2, -h/J_1]$ and $[-h/J_1, 1]$. On the lowest interval, $h + J_1 m > 0$ and $h + J_2 m < 0$, so for large β the value of the right hand side of (9 is approximately $p_1 - p_2$. In the middle interval both terms are positive, so for large β the right hand side is approximately $p_1 + p_2 = 1$. Finally, in the upper interval $h + J_1 m < 0$ and $h + J_2 m > 0$, so the value of the right hand side is approximately $p_2 - p_1$. Suppose $J_1 = -J_2$. When $p_1 = p_2$ there is only one steady state $m^* > 0$. But when p_1 is sufficiently small, so the the population is very biased towards positive conformity effects, two negative steady states emerge.

2.5. A general model with observable actions

In the last subsection we examined a variety of models in which actors had no information about the actions of others on which to condition their beliefs. Equilibrium in that system is the requirement that beliefs be rational, that is, that equilibrium beliefs be correct. We can also describe systems in which each actor has information about the realization of others' decisions. That is, $\mathcal{I}_i = (X_i, \omega_{-i})$. Choice probabilities for actor i now depend upon the choice of actor j. The specification of the model now describes conditional probabilities of the form

$$\text{Prob}(\omega_i|X, \omega_{-i}).$$

Any joint probability distribution consistent with these conditional distributions and the marginal distribution of the X_i's is an admissible characterization of the entire system. In other words, the models make assumptions about the nature of conditional distributions of a vector of random variables. Equilibrium in these models is a joint distribution consistent with the given conditional distributions.

An example of such a model is the Glaeser, Sacerdote and Scheinkman (1997) model of crime. The empirical models these issues present are discussed in Glaeser and Scheinkman (1998, this volume).

Suppose, for instance, that individuals can see the realization of choices in some subset N_i of the population. Equation (4) becomes

$$\text{Prob}(\omega_i = 1 | X_i, \omega_{-i}) = \frac{1}{1 + \exp -2\beta\Big(h(X_i) + \sum_{j \in N_i} E\big(J_{ij}(X_i, X_j)\omega_i\omega_j | X_i, \omega_{-i}\big)\Big)} \tag{10}$$

In the model of the previous section joint distributions given a value of Z were given by (5), and the equilibrium condition (3) — finding the right joint distribution of the ω_i's — was a matter of choosing the right value of Z. In this model the joint distribution no longer factors into an expression like equation (5), and one must search directly for the joint distribution.

Definition 2.1: A *Gibbs equilibrium* is a conditional distribution $\mu(\omega_1, \ldots, \omega_I | X_1, \ldots, X_I)$ such that the conditional distribution $\mu(\omega_i | X_i, \omega_{-i})$ is given by (10)

Notice to that the factorization of the joint distribution of choice into the product of the distributions of individual choice is lost because each individual choice depends upon the realization of the choices of others — rather than some estimate of the choices of others, as was the case previously. This model falls into a well-studied class of probabilistic systems called *random fields*. Random fields describe the joint distribution of a large set of random variables. Given are distributions for the realization of each random variable conditional on the others. The object of the study is to demonstrate the existence and characterize the joint distribution through the specification of the conditionals. The joint distributions are sometimes known as Gibbs measures due to their importance and historical origins in statistical mechanics — thus the name above, which we hope will not become conventional. The existence problem for Gibbs measures is straightforward for finite populations, but the large population asymptotic behavior is difficult. Below we will give some examples in which equilibrium fails to exist. See Georgii (1988).

Again for expositional and computational simplicity, we will assume for the remainder of this section that there is no uncertainty due to lack of knowledge about actors' characteristics. That is, for all i, $\mathcal{I}_i = (X, \omega_{-i})$.

2.5.1 Uniform global interaction

Suppose that actors are homogeneous (no $X_i's$) and

$$J_{i,j} = \begin{cases} J/(I-1) & \text{if } i \neq j, \\ 0 & \text{if } i = j. \end{cases}$$

Equation (10) becomes

$$\begin{aligned} \text{Prob}(\omega_i|\omega_{-i}) &= \frac{1}{1+\exp -2\beta\big(h\omega_i + \frac{J}{I-1}\sum_{j\neq i}\omega_i\omega_j\big)} \qquad (10) \\ &= \frac{\exp\beta\big(h\omega_i + \frac{J}{I-1}\sum_{j\neq i}\omega_i\omega_j\big)}{\exp\beta\big(h\omega_i + \frac{J}{I-1}\sum_{j\neq i}\omega_i\omega_j\big) + \exp -\beta\big(h\omega_i + \frac{J}{I-1}\sum_{j\neq i}\omega_i\omega_j\big)} \end{aligned}$$

A computation verifies that the Gibbs equilibrium distribution is

$$\text{Prob}(\omega) \sim \exp\beta\bigg(\sum_i h\omega_i + \sum_{i,j}\frac{J}{2(I-1)}\omega_i\omega_j\bigg) \qquad (12)$$

The function

$$H(\omega) = \sum_i h\omega_i + \sum_{i,j}\frac{J}{2(I-1)}\omega_i\omega_j$$

is known as the *potential function,* and it characterizes equilibrium.[4] Equation (12) is simply written as $\text{Prob}(\omega) \sim \exp\beta H(\omega)$. As the parameter β diverges, the equilibrium probability distribution concentrates its mass on the maxima of H. For instance, when $J > 0$, the distribution will concentrate on $\omega_i \equiv 1$, $\omega_i \equiv 0$ or probability (1/2) on each if h is positive, negative or 0, respectively.

It is interesting to compare the large β behavior of this example to that of section 1.2.1 when $h = 0$. In that example there were three solutions: 0 and $\pm m_\beta$, and $m_\beta \to 1$ as $\beta \to \infty$. In the example of this section, the mean is always 0, but the distribution is concentrating on the two configurations in which all players agree. The probability of agreement converges to 1 as β diverges. The configurations in which mean play is 0, are of minimal probability for all β, and that probability goes to 0 rapidly as β grows. So the $m = 0$ solution really does not appear here.

---

[4] In physical models, the potential is taken to be $-H$, and "maxima" below is replaced with "minima". Our sign convention is consistent with game theorists' discussion of potential games. Connections will be drawn below.

2.5.2 Uniform local interaction

The local interaction described in section 1.2.2 above can also be modeled with observable actions. Suppose again, as in that section, that all actors are homogeneous. Represent them as the vertices of a connected graph and let n_i denote the number of neighbors of i. Let $J_{i,j} = J$ if i and j are neighbors, and 0 otherwise. Equation (10) becomes

$$\begin{aligned}\operatorname{Prob}(\omega_i|\omega) &= \frac{1}{1+\exp -2\beta\big(h\omega_i + \frac{J}{n_i}\sum_{\{i,j\}\ni i}\omega_i\omega_j\big)} \\ &= \frac{\exp\beta\big(h\omega_i + \frac{J}{n_i}\sum_{\{i,j\}\ni i}\omega_i\omega_j\big)}{\exp\beta\big(h\omega_i + \frac{J}{n_i}\sum_{\{i,j\}\ni i}\omega_i\omega_j\big) + \exp -\beta\big(h\omega_i + \frac{J}{n_i}\sum_{\{i,j\}\ni i}\omega_i\omega_j\big)}\end{aligned} \tag{12}$$

where the sums indexed by $\{i, j\}$ are over all unordered pairs of vertices connected by an edge in the graph.

Imagine first that the graph is a circle. Each actor has two neighbors, so $n_i = 2$. The equilibrium is $\operatorname{Prob}(\omega) \sim \exp\beta H(\omega)$ where

$$H(\omega) = \sum_i h\omega_i + \sum_{\{i,i+1\}} \frac{J}{2}\omega_i\omega_j$$

where $I + 1$ is taken to be 1.

2.5.3 Non-existence of Gibbs equilibrium — local interaction

Next suppose that there are three actors on a line. Actors 1 and 3 have only one neighbor while actor 2 has two neighbors. For this case no potential exists. There is no probability distribution with conditional probabilities given by (13). To see this, suppose the contrary. Then there is a potential function H such that $\operatorname{Prob}(\omega) = \exp\beta H(\omega)$. (It should be clear that *any* completely mixed probability distribution on any finite state space can be written this way.) Consequently

$$\operatorname{Prob}(\omega_1 = 1|\omega_2, \omega_3) = \frac{1}{1+\exp\beta\big(H(-1,\omega_2,\omega_3) - H(1,\omega_2,\omega_3)\big)}$$

and similarly for the other two actors. Equation (11) has

$$\operatorname{Prob}(\omega_1 = 1|\omega_2, \omega_3) = \frac{1}{1+\exp -2\beta(h + J\omega_2)}.$$

Consequently

$$H(1,\omega_2,\omega_3) - H(-1,\omega_2,\omega_3) = 2(h + J\omega_2).$$

Similarly,

$$H(\omega_1, 1, \omega_3) - H(\omega_1, -1, \omega_3) = 2\big(h + \frac{J}{2}(\omega_1 + \omega_3)\big),$$
$$H(\omega_1, \omega_2, 1) - H(\omega_1, \omega_2, -1) = 2(h + J\omega_2).$$

The argument consists of showing that no function can have these first differences. To see this, observe that $H(1, 1, -1) - H(-1, -1, -1)$ can be written in two ways:

$$\begin{aligned} H(1,1,-1) - H(-1,-1,-1) =& H(1,1,-1) - H(1,-1,-1) + \\ & H(1,-1,-1) - H(-1,-1,-1) \\ =& H(1,1,-1) - H(-1,1,-1) + \\ & H(-1,1,-1) - H(-1,-1,-1). \end{aligned}$$

The value of the first path is

$$2h + 2(h - J) = 4h - 2J.$$

The value of the second path is

$$2(h + J) + 2(h - J) = 4h.$$

If and only if $J = 0$, the case of no interactions, are the two expressions the same. Consequently an equilibrium exists if and only if $J = 0$.

In fact it is easy to see that in any local interaction model with choice probabilities given by (13), equilibrium will exist if and only if for all i and j, $n_i = n_j$. Other specifications are possible. For instance, if we do not divide by n_i, then an equilibrium will exist for any graph. This has the effect, which may or may not be desirable, of assuming that people with more neighbors put more weight on conforming than those with fewer neighbors.

Notice that although no Gibbs equilibrium exists for this example, there is an equilibrium in the sense of section 2.2.

2.5.4 Non-existence of Gibbs equilibrium — global interaction

Generally speaking, the model of equation (10) can allow for a moderate degree of heterogeneity. Suppose that each actor has her own β_i, h_i and J_{ij} parameters, so that equation (10) is

$$\text{Prob}(\omega_i|\omega_{-i}) = \frac{1}{1 + \exp 2\beta_i(h_i\omega_i + \sum_{j\in N_i} J_{ij}\omega_i\omega_j)}$$

A computation shows that equilibrium exists if and only if $\beta_i J_{ij} = \beta_j J_{ji}$. The equilibrium distribution has

$$\rho(\omega) \sim \exp \sum_i \beta_i \left(h_i\omega_i + \frac{1}{2}\sum_j J_{ij}\omega_i\omega_j \right).$$

In general, the existence of equilibrium requires a certain symmetry of the interaction effects in order to guarantee the path independence illustrated in the example.

A special case of this calculation covers heterogeneous preferences and global interaction. In this case $N_i = \{1, \ldots, I\}/\{i\}$, $\beta_i = \beta_j$, and for each i there is a J_i such that for all j, $J_{ij} \equiv J_i$. The existence of Gibbs equilibria rests entirely on the social interaction effects, the J_i. Gibbs equilibria exist if and only if all actors are identical in the social interaction component of their preferences. This suggests that Gibbs equilibria is not a particularly useful equilibrium concept.

2.6. dynamics

The models of section 2.2 are game theoretic. Actors have beliefs about the choice of others, and they respond to their beliefs. The only difference between the equilibrium models of section 2.2 and conventional Nash equilibrium is the model of individual choice. Section 2.2 replaces expected utility optimization with random utility models in the spirit of Block and Marschak (1960).[5] On the other hand, conditioning directly on the value of observables as we do in section 2.3 is more in the spirit of rational expectations equilibrium. In REE models no actor is strategic with respect to the information revealed by her choice, presumably because of the large number of market participants and the smallness of each individual actor. But when actors are strategic existence issues arise. With deterministic choice we might easily see no equilibrium because of best-response cycles. If we see Gibbs equilibrium only because we have fuzzed up choice enough, these models would be rather unsatisfactory. But as unnatural as these models seem from the point of view of conventional models of strategic interaction, they are important objects of empirical study because they describe the stationary distributions of interesting and well-motivated dynamical processes in which, at random moments, actors playing a game with each other revise their current choices based on the current actions of their neighbors. These models were introduced by Blume (1993), Kandori, Mailath and Robb (1993) and Young (1993). The discussion here follows Blume (1993, 1994).

2.6.1 Uniform global interaction

Consider a population of I actors interacting with one another. At the beginning of time each actor is assigned an action. This assignment is the *initial configuration* of the system. Each actor has a rate-1 Poisson alarm clock independent of all others. When her alarm clock rings, she has an opportunity to revise her strategy. Put formally, each actor i is endowed with a collection of random variables $\{\tau^i_n\}_{n=1}^{\infty}$ such that each $\tau^i_n - \tau^i_{n-1}$ is exponentially distributed with mean 1, and all such differences are independent of all others, hers and the other actors'. When actor i has an opportunity to revise her action, her choice is described by the probability distribution (11). So actors respond myopically to the current behavior of the population.[6] The choices of the population matter for actor i only through the sample average $I^{-1}\sum_{j\neq i}\omega_j$. At a revision opportunity for actor i, the

---

[5] For a discussion of this view, see Blume, Holt and Salant (1989) and Blume (1993).

[6] For an example of non-myopic choice, see Blume (TK).

sample average of everyone but her can be recovered from knowledge of her choice and the aggregate choice of the population.

Let $\{S_t\}$ denote the the stochastic process whose value at time t is $S_t = \sum_j \omega_{it}$. This process changes state whenever an actor changes her choice. If an actor changes from a -1 to a $+1$ S_t increases by 2, and it decreases by 2 whenever an actor changes in the opposite direction. The transition rates can be computed from the conditional probability distribution (11). Suppose the system is in state S. It transits to state $S+2$ only when a revision opportunity comes to one of the $(S-I)/2$ actors currently choosing -1, and that actor chooses $+1$.[7] The probability of a -1 actor making this choice is

$$\frac{1}{1+\exp -2\beta\bigl(h+\frac{J}{I-1}(S+1)\bigr)}\,.$$

Consequently transition rate from S to $S+2$ is

$$\lambda_S = \frac{I-S}{2}\frac{1}{1+\exp -2\beta\bigl(h+\frac{J}{I-1}(S+1)\bigr)}\,.$$

A similar computation gives the transition rate in the other direction. To transit from $S+2$ back to S requires that one of the $(S+2+I)/2$ actors choosing $+1$ switches to -1. The transition rate is

$$\mu_{S+2} = \frac{I+S+2}{2}\frac{1}{1+\exp 2\beta\bigl(h+\frac{J}{I-1}(S-1)\bigr)}\,.$$

The process of sums is an example of a birth-death process. The states are $\{-I, -I+2, \dots, I-2, I\}$, and the invariant distribution is characterized by the conditions

$$\text{Prob}(S)\lambda_S = \text{Prob}(S+2)\mu_{S+2}\,.$$

Consequently

$$\begin{aligned}\frac{\text{Prob}(S+2)}{\text{Prob}(S)} &= \frac{\lambda_S}{\mu_{S+2}} \\ &= \frac{I-S}{I+S+2}\frac{\exp\beta\bigl(h+\frac{J}{I-1}(S+1)\bigr)}{\exp -\beta\bigl(h+\frac{J}{I-1}(S+1)\bigr)} \\ &= \exp 2\beta\bigl(h+\frac{J}{I-1}(S+1)\bigr)\,.\end{aligned}$$

Consequently,

[7] There are other imaginable transitions, such as where two -1 actors switch to +1 and one +1 actor switches to -1, but these transitions all involve the simultaneous arrival of revision opportunities to more than one actor, and is thus a 0-probability transition.

$$\text{Prob}(S) \sim \binom{I}{\frac{S+I}{2}} \exp\beta\big(hS + \frac{J}{2(I-1)}S^2\big)$$

All configurations summing to the same total S are equally likely, and their number is given by the binary coefficient in the equation, so for any given configuration $(\omega_1, \ldots, \omega_I)$:

$$\text{Prob}(\omega_1, \ldots, \omega_I) \sim \exp\beta\left(h\sum_i \omega_i + \frac{J}{2(I-1)}\sum_{\{i,j\}} \omega_i\omega_j\right). \tag{14}$$

The point of this exercise is that (14) and (12) are identical. The stationary distribution for the process (14) is the same as the Gibbs equilibrium (12). So even though the distribution described by (12) may not be appealing as a static equilibrium, it arises quite naturally as a description of long-run behavior of the process.

The short run behavior of the process $\{S_t\}_{t\geq 0}$ from any given state m and time t can be estimated by looking at the evolution of the conditional means $E(S_\tau | S_t = S)$. This estimate is a very good approximation of the path of the process over any finite time interval $[t, t+T]$ for large population sizes.

Let $s_t = S_t/I$. The process $\{s_t\}_{t\geq 0}$ records averages. Suppose the process is in state s at time t. Let f denote any function of the state variable s. The expected value of $f(s)$ at time $t+\tau$ can be computed from the transition rules. It is a solution to the differential equation

$$\begin{aligned}
\frac{dE\big(f(s_{t+\tau})|s_t = s\big)}{dt} &= \frac{I(1-s)}{2}\frac{1}{1+\exp -2\beta(h + J\frac{I}{I-1}s - \frac{J}{I-1})}\left(f\left(s+\frac{2}{I}\right) - f(s)\right) + \\
&\quad \frac{I(1+s)}{2}\frac{1}{1+\exp 2\beta(h + J\frac{I}{I-1}s - \frac{J}{I-1})}\left(f\left(s-\frac{2}{I}\right) - f(s)\right) \\
&= \frac{I(1-s)}{2}\frac{1}{1+\exp -2\beta(h + J\frac{I}{I-1}s - \frac{J}{I-1})}f'(s)\frac{2}{I} - \\
&\quad \frac{I(1+s)}{2}\frac{1}{1+\exp 2\beta(h + J\frac{I}{I-1}s - \frac{J}{I-1})}f'(s)\frac{2}{I} + Io(I) \\
&\stackrel{I\to\infty}{\longrightarrow} \big(\tanh\beta(h+Js) - s\big)f'(s)\,.
\end{aligned}$$

The first equation holds for any continuous function of the state, the second for differentiable functions, and the third follows from letting I become large. Applying this to the function $f(s) = s$ gives, for large I,

$$\dot{m} = \tanh\beta(h + Jm) - m \tag{15}$$

The importance of equation (15) is that if I is large enough, the solution to (15) is a good approximation to the behavior of the stochastic process of averages.

Theorem: Let $\{s_t^I\}_{t\geq 0}$ refer to the average process with population size I, and suppose $s_0 = m_0$. Then for every $t \geq 0$,

$$\lim_{I\to\infty} \sup_{\tau\leq t} |s_\tau^I - m_\tau| = 0 \qquad \text{a.s.}$$

This Theorem is a standard result in the theory of density dependent population processes. (An elementary proof is too long to be given here. A quick high-tech proof can be found in Chapter 11.2 of Ethier and Kurtz, 1986.)

The steady states of this equation are the equilibria of the Brock-Durlauf uniform global interaction model (see equation (6)). When there is a unique equilibrium, it is a stable rest point of equation (15). When there are three equilibria, the center equilibrium is unstable and the two extreme equilibria are stable. But while the Brock-Durlauf model gives a good picture of the short-run behavior of the average process, it does not give sharp predictions of the asymptotic analysis. So long as $|h/J| < 1$ the differential equation will have (for large β) two sinks, which are convergingto ± 1. But as β becomes large, the stationary distribution puts most of its mass on the global maximum of the potential function $H(\omega)$, which is $\omega_i \equiv +1$ if $h > 0$ and $\omega_i \equiv -1$ if h is negative. As β grows the probability of a revision opportunity resulting in a non-optimal choice goes to 0. Consequently, for $h \neq 0$ one of the sinks, $\omega_i \equiv \text{sgn}(h)$, gets mass converging to 1 as β becomes large. This state is said to be stochastically stable.

2.6.2 Uniform local interaction

In the previous example we could have computed equilibrium in a different way. For a general continuous time Markov chain on a finite state space E, suppose the rate of transition from state e to state f is given by $q(e, f)$. A probability distribution ρ on E is a stationary distribution if it satisfies the *balance conditions* if for all e,

$$\sum_f \rho(f)q(f,e) = \sum_f \rho(e)q(e,f)$$

It satisfies the *detailed balance conditions* if for all e and f,

$$\rho(f)q(f,e) = \rho(e)q(e,f).$$

Clearly any solution to the detailed balance conditions is stationary. And for the processes we study, a distribution ρ solves the detailed balance conditions if and only if it is an equilibrium distribution.

For instance, suppose again that all actors have homogeneous preferences and live on the vertices of a connected graph. Suppose too that each actor has an indentical number

n of neighbors. Equation (4) becomes:

$$\text{Prob}(\omega_i|\omega_{-i}) = \frac{1}{1 + \exp 2\beta(h\omega_i + \frac{J}{n}\sum_{j\in N_i}\omega_i\omega_j)}$$

For a given configuration ω let ω^i denote the configuration such that:

$$\omega^i_j = \begin{cases} -\omega_i & \text{if } j = 1, \\ \omega_j & \text{otherwise.} \end{cases}$$

Assuming independent rate 1 Poisson alarm clocks, the transition rates are:

$$q(\omega, \omega') = \begin{cases} \text{Prob}(-\omega_i|\omega_{-i}) & \text{if for some } i,\ \omega' = \omega^i, \\ 0 & \text{otherwise.} \end{cases}$$

Consequently the detailed balance conditions are that for all i and ω:

$$\rho(\omega^i)q(\omega^i, \omega) = \rho(\omega)q(\omega, \omega^i)$$

The solution is:

$$\rho(\omega) \sim \exp\beta\Big(h\sum_i \omega_i + \frac{J}{2n}\sum_i\sum_{j\in N_i}\omega_i\omega_j\Big)$$

2.6.3 Actor heterogeneity — social valuations

When agents are homogeneous, Gibbs equilibria exist and are the stationary distributions of the Markov processes which describe population behavior. This need not be the case if actors differ in their characteristics. Consider the following global interaction model. There are N actors, Each actor is of one of two types. These types differ only with respect to the weight they put on the choices of others, J_1 or J_2.

Suppose that actors $1, \ldots, I_1$ are of type 1 and actors $I_1 + 1, \ldots, I_1 + I_2$ are of type 2. Let $\{S_{1t}, S_{2t}\}_{t\geq 0}$ denote the stochastic process whose value at time t is $S_{1t} = \sum_{i=1}^{I_1}\omega_{it}$ and $S_{2t} = \sum_{i=I_1+1}^{I_2}\omega_{it}$. The transition rules are defined in a manner analogous to that for the uniform global interaction process of section 2.4.1. The possible transitions from (S_1, S_2) are to $(S_1 \pm 2, S_2)$ or $(S_1, S_2 \pm 2)$. The transition rates can be computed from the conditional probability distribution (11).

Just as with uniform local interaction (sec. 2.4.1) the short run behavior for large I can be described by a differential equation — in this case, a pair of them. Let s_{kt} equal $(1/I)$ times the sum of the ω_i over the type k actors. Thus s_{kt} is related to the conditional mean m_{kt} by the relationship $s_{kt} = p_k m_{kt}$ where p_k is the fraction of the population consisting of type k individuals. Then for large I, the behavior of the system is approximately described by the equation system

$$\dot{s}_{kt} = p_k \tanh\beta\big(h + J_k(s_1 + s_2)\big) - s_{kt} \tag{16}$$

Let $\{s^I_{1t}, s^I_{2t}\}_{t\geq 0}$ refer to the process of the s_{kt} variables with population size I.

Theorem: Suppose $s^I_{k0} = s_{k0}$ for all I and $k = 1, 2$. Let s_{kt} denote a solution to (16) with initial conditions s_{k0}. Then for every $t \geq 0$,

$$\lim_{I\to\infty} \sup_{\tau\leq t} |s^I_{k\tau} - s_{k\tau}| = 0 \qquad \text{a.s.}$$

Again notice that the steady states are precisely the equilibria in the model of section 2.2.4.

Because this model has no hamiltonian, the detailed balance conditions have no solution. Invariant distributions are the solutions to the more general balance conditions. We hope that an examination of the balance conditions will clarify the relationships between the stationary distribution of the dynamic models and the equilibria of the Brock-Durlauf model.

2.7. Remarks

2.7.1 Empirical implications

The existence of an equilibrium distribution has two empirical consequences. The first is observable at the micro level. This is the absence of long range dependence. In the empirical distribution constructed from time-series observations on the local interaction model of section 2.4.2, we would see for large sample sizes that:

$$\text{pr}(\omega_i|\omega_{-i}) \approx \text{pr}(\omega_i|\omega_{N_i})$$

where ω_{N_i} is the projection of ω onto those coefficients which are in the set N_i. The long run behavior of actor i is independent of the choices of actors $j \notin N_i$. The appearance of long-range effects in the empirical distribution of a local interaction model is inconsistent with the existence of an equilibrium. In global interaction models, non-uniformities in the specification of the random utility model are reflected in the empirical distribution of states — for instance, if J_{ij} is some function of $|i - j|$, this relationship can be recovered from the empirical distribution.

The second empirical consequence is more subtle, but indicates the special nature of these processes. Suppose that $\{\omega_t\}_{t\geq 0}$ is a stationary process of configurations. Choose a t_0 and $0 < k < t_0$. Then the processes $\{\omega_t\}_{t_0-k<t<t_0+k}$ and $\{\omega_{-t}\}_{t_0-k<t<t_0+k}$ have the same joint distributions. This property is called *reversibility*, and is equivalent to the stationary distribution solving the detailed balance conditions rather than just the balance conditions. So, for instance, all stationary single-type birth-death processes are reversible.

2.7.2 Equilibrium, dynamics and game theory

The parameter β governs the sensitivity of choice to payoffs. If β is large each actor will best respond at a revision opportunity with high probability. The effect of large β on the stationary distribution is to concentrate most mass on the global maxima of the potential function. To understand the significance of this fact, suppose the social interaction in this system were to be modeled as a game. For the purposes of this discussion, suppose we are considering uniform global interaction. Each player has two strategies, $+1$ and -1, and the payoff of any configuration to player i is

$$h\omega_i + \frac{J}{I-1}\sum_j \omega_i\omega_j \tag{17}$$

If $J > 0$ this is a model of a coordination game in which each player has a preference for strategy $+1$. One interpretation of this model has been offered by Kandori, Mailath and Robb (1993) and Young (1993). Players are repeatedly matched randomly (and uniformly so) against other players in the population for pairwise interactions. In other words, players are randomly matched to play a symmetric two-person game. Equation (17) represents the return against a randomly drawn opponent.[8] In this case the two-person game is a coordination game if $J > 0$, and $h > 0$ if and only if the strategy $+1$ is risk-dominant. If $|h|$ is not too large, this game will have three equilibria: All play down, all play up, and all mix such that the expected value of $J\omega_j = -h$. These three equilibria correspond to the critical points of the potential function. The two pure strategy equilibria are local maxima, while the mixed strategy is the global minimum.

The dynamic process of section 2.4 can be thought of as a process of playing the game. At random moments individuals have an opportunity to revise their strategy. Their payoffs are given by (17). But individuals are boundedly rational, and occasionally make mistakes, so the revision process is in fact stochastic with strategy selection probabilities given by (11). Long run play of the game is described by the invariant distribution. Much has been made of the behavior of this system for large β. Kandori, Mailath and Robb (1993) and Young (1993), and Blume (1993) for local interaction, observe that the distribution will concentrate around $\omega_i \equiv +1$ or $\omega_i \equiv -1$ depending upon the sign of h. In the language of Foster and Young (TK), the risk-dominant equilibrium is *stochastically stable.*

As we mentioned above, the model of section 2.2 also has a game theory interpretation. It is the natural extension to a random utility choice framework of the conventional expected utility maximization paradigm with the payoff function (14). Models like this have been investigated by Blume, Holt and Salant (1989) and McKelvey and Palfrey (1996). Notice that for coordination games, as β grows the equilibria of the random choice models converge to the Nash equilibrium set. Upper hemi-continuity at the $\beta \to +\infty$ limit, but it is probable that lower hemi-continuity fails in games with Nash equilibria in weakly dominated strategies.

---

[8] They did not put the model this way. The equivalence of the two can be seen by rewriting the payoff matrix as in Blume (1993). They also used a different random utility model. We discuss this below.

Of course this analysis applies to a much larger class of games than two-by-two coordination games. The general class of I-person interactions to which this method (with extreme-value random utility choice) applies is the class of I-player *potential games,* which were introduced by Monderer and Shapley (1996). Briefly, a potential is a function H on the space of configurations such that $H(s, \omega_{-i}) - H(s', \omega_{-i})$ gives the payoff gain to player i of switching fromstrategy s to strategy s'. Any local maximum of the potential function is a Nash equilibrium. When the game is extended to mixed strategies, any interior critical point is also a Nash equilibrium. The class of potential games is mathematically small but economically important. Monderer and Shapley suggest that, in selecting among equilibria, a privileged place might go to the global maximum of the potential function. The dynamic analysis provides a justification for this equilibrium refinement. For games without a potential in the sense of Monderer and Shapley, the equilibria of section 2.3 will fail to exist.

2.7.3 Other random utility models

Despite the strangeness of the extreme value distribution, the specifications (4) and (10) for individual choice behavior are quite natural. They say simply that the log-odds of choosing $+1$ over -1 is proportional to their payoff difference — a natural logit specification. Kandori, Mailath and Robb (1993) and Young (1993) use a different model of randomness in choice. They assume that agents best respond with some probability $1 - \epsilon$, and make a mistake with probability ϵ. Blume (1994) and Maruta (1996) have studied the extent to which this analysis can be carried out with uniform global interaction. Blume considered stochastic binary choice in which the log-odds of choosing $+1$ over -1 is proportional to a function g of the payoff difference Δ, and showed that all g skew-symmetric functions g generate stationary distributions with the same large β behavior. This include the Kandori, Mailath and Robb and Young choice models, where $g(\Delta) = \mathrm{sgn}(\Delta)$. Among other things, Maruta (1996) demonstrated that skew-symmetry was necessary as well as sufficient to ensure that the risk-dominant configuration is stochastically stable in all games. Thus skew-symmetry characterizes the degree to which the large β behavior of equilibrium is robust to the specification of choice probabilities. Blume (1994) also gives a necessary and sufficient condition on the distribution of ϵ in equation $(4')$ for the same characterization to work — namely, that conditional upon being in the tails of the distribution, neither the upper nor the lower tail is infinitely more likely than the other. These results to not generalize to local interaction models. With the "mistakes" model, the equilibrium of section 2.3 will fail to exist and the set of stochastically stable states need not be that of the extreme value rule.