Social Interactions: Theory

Steven N. Durlauf

University of Wisconsin at Madison

Complementaries

Behind social interactions models is the assumption that complementarities exist between the behavior of individuals. This idea has been very extensively explored in the economic literature, perhaps most deeply in the work on Paul Milgrom and John Roberts.

Social interactions models are typically much less sophisticated than those studied in the game theory literature (although there are exceptions!!)

Example: Cooper and John

Cooper and John's (1988) paper illustrates the main ideas in modeling complementarities among economic agents.

In their model, they consider I agents, each of whom makes an effort choice $e_i \in [0,1]$.

Each agent has a payoff function

$$V\left(e_i, \bar{e}_{-i}\right)$$

where $\bar{e}_{-i} = \left(I-1\right)^{-1} \sum_{j \neq i} e_j$

The payoff function is assumed twice differentiable. Comment: I will not worry about corner solutions in the discussion.

The key to the Cooper and John analysis is the assumption that the payoff function exhibits complementarities.

$$\frac{\partial^2 V\left(e_i, \overline{e}_{-i}\right)}{\partial e_i \partial \overline{e}_{-i}} > 0$$

Note this assumption means that for effort levels $a > b$ and $c > d$,

$$\int_b^a \int_d^c \frac{\partial^2 V\left(e_i, \overline{e}_{-i}\right)}{\partial e_i \partial \overline{e}_{-i}}\, de_i d\overline{e}_{-i} =$$

$$V\left(a, c\right) - V\left(b, c\right) - \left(V\left(a, d\right) - V\left(b, d\right)\right) > 0$$

which can be rewritten

$$V\left(a, \overline{e}_{-i}\right) - V\left(b, \overline{e}_{-i}\right) > V\left(a, \underline{e}_{-i}\right) - V\left(b, \underline{e}_{-i}\right)$$

Critical Idea

Complementarities induce a tendency towards similar behavior.

Equilibria

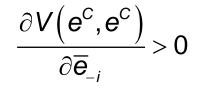A symmetric Nash equilibrium is an effort level e^{NC} such that

$$\frac{\partial V\left(e^{NC},e^{NC}\right)}{\partial e_i}=0$$

In contrast, a cooperative equilibrium is an effort level e^{C} such that

$$\frac{\partial V\left(e^{C},e^{C}\right)}{\partial e_i}+\frac{\partial V\left(e^{C},e^{C}\right)}{\partial \overline{e}_{-i}}=0$$

So the cooperative and noncooperative equilibria will not coincide unless

$$\frac{\partial V\left(e^{c}, e^{c}\right)}{\partial \overline{e}_{-i}} = 0$$

If

$$\frac{\partial V\left(e^{C}, e^{C}\right)}{\partial \bar{e}_{-i}} > 0$$

then the noncooperative equilibrium implies socially inefficient effort.

Comment: Milgrom and Robert extend to vector choices, payoffs with discontinuities, noncontinuous choice spaces.

Statistical Mechanics

Statistical mechanics is a branch of physics which studies the aggregate behavior of large populations of objects, typically atoms.

A canonical question in statistical mechanics is how magnets can appear in nature. A magnet is a piece of iron with the property that atoms tend on average to be spinning up or down; the greater the lopsidedness the stronger the magnet. (Spin is binary).

While one explanation would be that there is simply a tendency for individual atoms to spin one way versus another, the remarkable finding in the physics literature is that interdependences in spin probabilities between the atoms can, when strong enough, themselves be a source of magnetization.

Classic structures of this type include the Ising and Currie-Weiss models.

Economists of course have no interest in the physics of such systems. On the other hand, the mathematics of statistical mechanics has proven to be useful for a number of modeling contexts. As illustrated by the magnetism example, statistical mechanics models provide a language for modeling interacting populations.

The mathematical models of statistical mechanics are sometimes called interacting particle systems or random fields, where the latter term refers to interdependent populations with arbitrary index sets, as opposed to a variables indexed by time.

Statistical mechanics models are useful to economists as these methods provide a framework for linking microeconomic specifications to macroeconomic outcomes.

A key feature of a statistical mechanical system is that even though the individual elements may be unpredictable, order appears at an aggregate level.

At one level, this is an unsurprising property; laws of large numbers provide a similar linkage. However, in statistical mechanics models, properties can emerge at an aggregate level that are not describable at the individual level.

Magnetism is one example of this as it is a feature of a system not an individual element.

The existence of aggregate properties without individual analogues is sometimes known as emergence.

As such, emergence is a way, in light of Sonnenschein-type results on the lack of empirical implications to general equilibrium theory, to make progress on understanding aggregate behavior in the presence of heterogeneous agents

The general structure of statistical mechanics models may be understood as follows.

Consider a population of elements ω_a, where *a* is an element of some arbitrary index set *A*.

Let ω denote vector all elements in the population and ω_{-a} denote all the elements of the population other than *a*.

Concretely, each ω_i may be thought of as an individual choice.

A statistical mechanics model is specified by the set of probability measures

$$\mu\left(\omega_a \middle| \omega_{-a}\right) \qquad (1)$$

for all *i.* These probability measures describe how each element of a system behaves given the behavior of other elements.

The objective of the analysis of the system is to understand the joint probability measures for the entire system,

$$\mu(\omega) \qquad\qquad (2)$$

that are compatible with the conditional probability measures.

Thus, the goal of the exercise is to understand the probability measure for the population of choices given the conditional decision structure for each choice. Stated this way, one can see how statistical mechanics models are conceptually similar to various game-theory models, an idea found in Blume (1993).

Dynamic versions of statistical mechanics models are usually modeled in continuous time. One considers the process $\omega_i(t)$ and unlike the atemporal case, probabilities are assigned to at each point in time to the probability of a change in the current value.

Operationally, this means that for sufficiently small δ

$$\mu\left(\omega_a(t+\delta)\big|\omega_a(t+\delta) \neq \omega_a(t)\right) = f\left(\omega_{-a}(t), \omega_a(t)\right)\delta + o(\delta) \quad (3)$$

What this means is that at each t, there is a small probability that $\omega_i(t)$ will change value, such a change is known as a flip when the support of $\omega_a(t)$ is binary. This probability is modeled as depending on the current value of element a as well as on the current (time t) configuration of the rest of the population. Since time is continuous whereas the index set is countable, the probability that two elements change at the same time is 0 when the change probabilities are independent.

Systems of this type lead to question of the existence and nature of invariant or limiting probability measures for the population, i.e. the study of

$$\lim_{t \Rightarrow \infty} \mu\big(\omega(t)\big|\omega(0)\big) \qquad\qquad (4)$$

Discrete time systems can of course be defined analogously; for such systems a typical element is $\omega_{a,t}$.

Important Caveat

This formulation of statistical mechanics models, with conditional probability measures representing the micro-level description of the system, and associated joint probability measures the macro-level or equilibrium description of the system, also illustrates an important difference between physics and economics reasoning.

For the physicist, treating conditional probability measures as primitive objects in modeling is natural. One does not ask "why" one atom's behavior reacts to other atoms. In contrast, conditional probabilities are not natural modeling primitives to an economics.

A Math Trick

The conditional probability structure described by (1) can lead to very complicated calculations for the joint probabilities (2). In the interests of analytical tractability, physicists have developed a set of methods referred to as mean field analyses.

These methods typically involve replacing the conditioning elements in (1) with their expected values, i.e.

$$\mu\left(\omega_a \big| E\left(\omega_{-a}\right)\right) \qquad (5)$$

A range of results exist on how mean field approximation relate to the original probabilities models they approximate.

From the perspective of economic reasoning, mean field approximations have a substantive economic interpretation as they implicitly mean that agents make decisions based on their beliefs about the behaviors of others rather than the behaviors themselves.

Markov Random Fields

An important class of statistical mechanics models generalizes the Markov property of time series to general index sets.

Definition 1. Neighborhood.

Let $a \in A$. A neighborhood of a , N_a is defined as a collection of indices such that

  i. $a \neq N_a$

  ii. $a \in N_b \Leftrightarrow b \in N_a$

    Neighborhoods can overlap.

The collection of individual neighborhoods provides generalization of the notion of a Markov process to more general index sets than time.

Definition 2. Markov random field.

Given a set of neighborhoods N_a, if $\forall a$

$$\mu(\omega) \Rightarrow$$
$$\mu\left(\omega_a \middle| \omega_{-a}\right) = \mu\left(\omega_a \middle| \omega_b \forall b \in N_a\right) \qquad (2.6)$$

then $\mu(\omega)$ is a Markov random field with respect to the neighborhood system.

For binary variables, again coded -1 and 1, there are some well known examples of random fields on Z^d

Hammersley-Clifford Theorem

The Hammersley-Clifford theorem provides a general functional form for the joint probability measure of a Markov random field when the support of the individual random variables is binary. In order to describe the theorem, it is necessary to introduce an additional type of subset of the indices, called a clique.

Definition 5. Clique.

Given a neighborhood collection N_a, a subset of the indices c is a clique if each member of the set is a neighbor of each of the other members.

With the definition of a clique, one can state the Hammersley-Clifford theorem.

Hammersley-Clifford theorem.

$\mu(\omega)$ is the probability measure of a Markov random field if and only if

$$\mu(\omega) = \exp\left(\sum_{c \in C} V_c(\omega)\right) \qquad (2.7)$$

where C denotes a collection of cliques and the value of $V_c(\omega)$ only depends on those elements of ω whose indices lie in c.

Some additional definitions are standard in the statistical mechanics literature.

Definition 6. Potential.

Any set of functions $V_a(\omega)$ defined over all subsets a of the indices is called a potential.

Definition 7. Neighborhood potential.

A potential is a neighborhood potential if $V_a(\omega) = 0$ is not a clique.

Definition 8. Gibbs measure.

The functional form $\exp \sum_{a \in A} V_a(\omega)$ is called a Gibbs measure.

Definition 9. Neighborhood Gibbs measure.

The functional form $\exp\left(\sum_{c \in C} V_c(\omega)\right)$ is called a neighborhood Gibbs measure.

When suitably normalized, the Gibbs measures can represent probability measures. Thus, the Hammersley-Clifford theorem states that all Markov random fields can be represented by neighborhood Gibbs probability measures.

The basic idea of the proof is to consider

$$\log \mu(\omega) - \log(-\underline{1}) = G(\omega) \qquad (2.8)$$

$\log \mu(-\underline{1})$ denotes the probability of a configuration in which each element equals -1. Moving from finite to countably infinite index sets is not a problem.

$G(\omega)$ is expanded as

$$G(\omega) =$$

$$\sum_a g_i(\omega_a) + \sum_a \sum_{b \neq a} g_{ab}(\omega_a \omega_b) \omega_a \omega_b + \sum_a \sum_{b \neq a} \sum_{c \neq a,b} g_{abc}(\omega_a \omega_b \omega_c) \omega_a \omega_b \omega_c + \ldots$$

$$(2.9)$$

The proof involves showing that all g functions that are not associated with cliques equal zero. See Cressie (1993) for details.

Properties of Statistical Mechanics Models

i. existence

The first question one naturally asks for environments of the type described concerns the existence of joint or invariant probability measure over a population of elements in which conditional probabilities for the behaviors of the elements have been specified.

Existence results of this type differ from classic results such as the Kolmogorov extension theorem in that they concern the relationship between conditional probabilities and joint probabilities, rather than relationship (as occurs in the Kolmogorov case) between joint probabilities measure on finite sets of elements versus an infinite collection that represents the union of the various elements.

Existence theorems are quite technical but do not, in my judgment, require conditions that are implausible from the perspective of socioeconomic systems.

ii. uniqueness or multiplicity

The existence of a joint or invariant measure says nothing about how many such measures exist.

When there are multiple measures compatible with the conditional probabilities, the system is said to be nonergodic.

Notice that for the dynamical models, the uniqueness question involves the dependence of the invariant measure on the initial configuration on $\omega(0)$ or ω_0.

Heuristically, for atemporal models, nonergodicity is thus the probabilistic analog to multiple equilibria whereas for temporal models, nonergodicity is the probabilistic analog to multiple steady states.

One of the fascinating features of statistical mechanics models is their capacity to exhibit nonergodicity in nontrivial cases.

Specifically, nonergodicity can occur when the various direct and indirect connections between individuals in a population create sufficient aggregate independence across agents.

As such statistical mechanics models use richer interactions structure than appear, for example in conventional time series model.

To see this, consider a Markov chain where

$$\Pr\left(\omega_t = 1 \middle| \omega_{t-1} = 1\right) \neq 1 \text{ and } \Pr\left(\omega_t = 0 \middle| \omega_{t-1} = 0\right) \neq 1$$

For this case

$\lim_{j \Rightarrow \infty} \Pr\left(\omega_{t+j} \middle| \omega_0\right)$ will not depend on ω_0.

However, suppose that $I = Z^0 \times Z$, i.e. the index set is the Cartesian product of the non-negative integers with the integers. I use i to capture denote the fact that its support is a cross product of integer lattices. Suppose that the system has a local Markov property of the form

$$\Pr\left(\omega_{i,t} \big| \underset{\sim}{\omega}_{t-1}\right) = \Pr\left(\omega_{i,t} \big| \omega_{i-1,t-1}\omega_{i,t-1}\omega_{i+1,t-1}\right)$$

in words, the behavior of a particular $\omega_{i,t}$ depends on its value at $t-1$ as well as its "nearest neighbors."

In this case, it is possible that $\lim_{j \Rightarrow \infty} \Pr\left(\omega_{i,t+j} \big| \underset{\sim}{\omega}_0\right)$ does depend on $\underset{\sim}{\omega}_0$ even though no conditional probability $\Pr\left(\omega_{i,t} \big| \omega_{i-1,t-1}\omega_{i,t-1}\omega_{i+1,t-1}\right)$ equals 1.

The reason for this is that in the case of an evolving set of Markov processes, there are many indirect connections.

For example, the realization of $\omega_{i-2,t-2}$ will affect $\omega_{i,t}$ because of its effect on $\omega_{i-1,t-1}$; no analogous property exists when there is a single element at each point in time.

In fact, the number of elements at time $t-k$ that affect $\omega_{i,t}$ is, in this example, growing in k.

This does mean that such a system necessarily has multiple invariant measures, merely that it can when there is sufficient sensitivity of $\Pr\left(\omega_{i,t}\big|\underset{\sim}{\omega}_{t-1}\right) = \Pr\left(\omega_{i,t}\big|\omega_{i-1,t},\omega_{i,t-1},\omega_{i+1,t-1}\right)$ to the realizations of $\omega_{i-1,t}\omega_{i,t-1}$ and $\omega_{i+1,t-1}$.

For many statistical mechanics models, this dependence can be reduced to a single parameter.

For example the Ising model may be written

$$\Pr\left(\omega_i \,\middle|\, \omega_{-i}\right) \propto \exp\left(J \sum_{|i-j|=1} \omega_j \right);$$

so J fully characterizes the degree of interdependence.

In one dimension, the model is always ergodic, outside of trivial cases whereas for 2 dimensions it may not be.

In statistical mechanics model, one often finds threshold effects, i.e. when J is below some $J < \overline{J}$, the system exhibits a unique invariant measure, whereas if $J > \overline{J}$, multiple measures exist.

Another statistical mechanics model is the Curie Weiss model:

$$\Pr\left(\omega_i \middle| \omega_{-i}\right) \propto \exp\left(J\bar{\omega}_{-i}\right);$$

Where $\bar{\omega}_{-i}$ is the average of the system elements other than i. (And yes, I am skipping technical details since there are an infinite number of elements).

The mean field approximation for this model is

$$\Pr\left(\omega_i \middle| \omega_{-i}\right) \propto \exp\left(JE\omega\right)$$

Spin Glasses

Some statistical mechanics models are based on general formulations of the form

$$\mu\left(\omega_i \middle| \omega_{-i}\right) \propto \exp\left(\sum_{i \neq j} J_{ij} \omega_i \omega_j\right) \qquad (10)$$

In physics. J_{ij} is usually treated as a random variable. When it can take on positive and negative values, this system is called a spin glass.

Spin glasses can exhibit "frustration" which means that interactions can be conflicting.

Modeling Social Interactions

We consider I individuals who are members of a common group g. Our objective is to probabilistically describe the individual choices of each i, ω_i (a choice that is taken from the elements of some set of possible behaviors Ω_i) and thereby characterize the vector of choices of all members of the group, ω.

From the perspective of theoretical modeling, it is useful to distinguish between three sorts of influences on individual choices. These influences have different implications for how one models the choice problem.

These components are

h_i, a vector of deterministic (to the modeler) individual-specific characteristics associated with individual i,

ε_i, a vector of random individual-specific characteristics associated with i, i.i.d. across agents,

and

$\mu_i^e(\omega)$, the subjective beliefs individual i possesses about behaviors in the group, expressed as a probability measure over those behaviors.

Individual choices ω_i are characterized as representing the maximization of some payoff function V,

$$\omega_i = \text{argmax}_{\lambda \in \Omega_i} V\left(\lambda, h_i, \mu_i^e(\omega), \varepsilon_i\right) \qquad (11)$$

The decision problem facing an individual as a function of preferences (embodied in the specification of V), constraints (embodied in the specification of Ω_i) and beliefs (embodied in the specification of $\mu_i^e(\omega)$).

As such, the analysis is based on completely standard microeconomic reasoning to describe individual decisions.

Beliefs

This basic choice model is closed by imposing self-consistency between subjective beliefs $\mu_i^e(\omega)$ and the objective conditional probabilities $\mu(\omega \mid F_i)$, where F_i denotes the information available to agent i. We assume that agents know the deterministic characteristics of others as well as themselves and also understand the structure of the individual choice problems that are being solved.

This means that subjective beliefs obey

$$\mu_i^e(\omega) = \mu\left(\omega \middle| h_j, \mu_j^e(\omega) \ \forall j\right) \qquad (12)$$

where the right hand of this equation is the objective conditional probability measure generated by the model; self-consistency is equivalent to rational expectations in the usual sense.

From the perspective of modeling individual behaviors, it is typically assumed that agents do not account for the effect of their choices on the decisions of others via expectations formation.

In this sense, this framework employs a Bayes-Nash equilibrium concept.

A Multinomial Logit Approach to Social Interactions

1. Each agent faces a common choice set with L discrete possibilities, i.e. $\Omega_i = \{0, 1, \ldots, L-1\}$.

2. Each choice l produces a payoff for i according to:

$$V_{i,l} = h_{i,l} + Jp^e_{i,l} + \varepsilon_{i,l} \qquad (13)$$

3. Random utility terms $\varepsilon_{i,l}$ are independent across i and l and are doubly exponentially distributed with index parameter β,

$$\mu\left(\varepsilon_{i,l} \leq \varsigma\right) = \exp\left(-\exp\left(-\beta\varsigma + \gamma\right)\right) \qquad (14)$$

where γ is Euler's constant.

Characterizing Choices

These assumptions may be combined to produce a full description of the choice probabilities for each individual.

$$\mu\left(\omega_i = l \middle| h_{i,j}, p^e_{i,j} \; \forall j\right) =$$

$$\mu\left(\mathrm{argmax}_{j \in \{0...L-1\}} \; h_{i,j} + Jp^e_{i,j} + \varepsilon_{i,j} = l \middle| h_{i,j}, p^e_{i,j} \; \forall j\right) \quad (15)$$

The double exponential assumption for the random payoff terms leads to the canonical multinomial logit probability structure

$$\mu\left(\omega_i = I \middle| h_{i,j}, p^e_{i,j} \ \forall j\right) = \frac{\exp\left(\beta h_{i,I} + \beta J p^e_{i,I}\right)}{\displaystyle\sum_{j=0}^{L-1} \exp\left(\beta h_{i,j} + \beta J p^e_{i,j}\right)} \quad (16)$$

S

o the joint probabilities for all choices may be written as

$$\mu\left(\omega_1 = I_1, \ldots, \omega_I = I_I \middle| h_{i,j}, p^e_{i,j} \ \forall i, j\right) =$$

$$\prod_i \frac{\exp\left(\beta h_{i,I_i} + \beta J p^e_{i,I_i}\right)}{\displaystyle\sum_{j=0}^{L-1} \exp\left(\beta h_{i,j} + \beta J p^e_{i,j}\right)} \quad (17)$$

Self-Consistency of Beliefs

Self-consistent beliefs imply that the subjective choice probabilities p_l^e equal the objective expected values of the percentage of agents in the group who choose l, p_l, the structure of the model implies that

$$p_{i,l}^e = p_l = \int \frac{\exp\left(\beta h_{i,l} + \beta J p_l\right)}{\sum_{j=0}^{L-1} \exp\left(\beta h_{i,j} + \beta J p_j\right)} dF_h \qquad (18)$$

where F_h is the empirical probability distribution for the vector of deterministic terms $h_{i,l}$.

It is straightforward to verify that under the Brouwer fixed point theorem, at least one such fixed point exists, so this model always has at least one equilibrium set of self-consistent aggregate choice probabilities.

Characterizing Equilibria

To understand the properties of this model, it is useful to focus on the special case where $h_{i,l} = 0 \; \forall i, l$. For this special case, the choice probabilities (and hence the expected distribution of choices within a group) are completely determined by the compound parameter βJ.

An important question is whether and how the presence of interdependencies produces multiple equilibria for the choice probabilities in a neighborhood.

In order to develop some intuition as to why the number of equilibria is connected to the magnitude of βJ, it is helpful to consider two extreme cases for the compound parameter, namely $\beta J = 0$ and $\beta J = \infty$.

For the case $\beta J = 0$, one can immediately verify that there exists a unique equilibrium for the aggregate choice probabilities such that $p_l = \dfrac{1}{L}$ $\forall l$. This follows from the fact that under the assumption that all individual heterogeneity in choices come from the realizations of $\varepsilon_{i,l}$, a process whose elements are independent and identically distributed across choices and individuals. Since all agents are ex ante identical, the aggregate choice probabilities must be equal.

The case $\beta J = \infty$ is more complicated. The set of aggregate choice probabilities $p_l = \dfrac{1}{L}$ is also an equilibrium if $\beta J = \infty$ since conditional on these probabilities, the symmetries in payoffs associated with each choice that led to this equilibrium when $\beta J = 0$ are preserved as there is no difference in the social component of payoffs across choices.

However, this is not the only equilibrium. To see why this is so, observe that for any pair of choices l and l' for which the aggregate choice probabilities are nonzero, it must be the case that

$$\frac{p_l}{p_{l'}} = \frac{\exp\left(\beta J p_l\right)}{\exp\left(\beta J p_{l'}\right)} \qquad (19)$$

for any βJ. This follows from the fact that each agent is ex ante identical. Thus, it is immediate that any set of equilibrium probabilities that are bounded away from 0 will become equal as $\beta J \Rightarrow \infty$.

This condition is necessary as well as sufficient, so any configuration such that $p_l = \dfrac{1}{b}$ for some subset of b choices and $p_l = 0$ for the other $L - b$ choices is an equilibrium. Hence, for the case where $J = \infty$, there exist

$$\sum_{b=1}^{L} \binom{L}{b} = 2^L - 1$$

different equilibrium probability configurations. Recalling that β indexes the density of random utility and J measures the strength of interdependence between decisions, this case, when contrasted with $\beta J = 0$ illustrates why the strength of these interdependences and the

degree of heterogeneity in random utility interact to determine the number of equilibria.

These extreme cases may be refined to produce a more precise characterization of the relationship between the number of equilibria and the value of βJ.

Theorem 1. Multiple equilibria in the multinomial logit model with social interactions

Suppose that individual choices are characterized by eq. (12) with self-consistent beliefs, i.e., that beliefs are consistent with eq. (14)Assume

that $h_{i,l} = k \ \forall i, l$. Then there will exist at least three self-consistent choice

probabilities if $\dfrac{\beta J}{L} > 1$.

When $L = 2$, this theorem reduces to the characterization of multiple equilibria with binary choices in Brock and Durlauf (2001a).

I will exposit this model for comparison; note that the support of the choices is -1,1

For the binary choice model, self-consistency means that

$$m_{i,g}^e = m_g = 2\int F_\varepsilon \left(k + cX + dY_g + Jm_g \right) dF_{X|g} - 1. \quad (20)$$

where recall that $F_{X|g}$ is the empirical within-group distribution of X. The description of a process for individual choices combined with its associated self-consistency condition fully specifies a model.

Precise results may be obtained if one specifies the functional forms for $F_{X,g}$ and F_ε. For the binary case,

$$F_\varepsilon(z) = \frac{1}{1 + \exp(-z)} \quad\quad (21)$$

so that the model errors are as before negative exponentially distributed, and that $k + cX_i + dY_g = h$, so that this component of the payoff differential between the two choices is constant across group members.

For this special case,

$$\Pr\left(\omega_i = 1 \middle| X_i, Y_g, g\right) = \frac{\exp\left(h + Jm_{i,g}^e\right)}{\exp\left(h + Jm_{i,g}^e\right) + \exp\left(-h - Jm_{i,g}^e\right)} \quad (22)$$

Under self-consistency, the expected average choice level m_g within a group must obey

$$m_g = \tanh\left(h + Jm_g\right) \qquad\qquad (23)$$

In (19), $\tanh(x) = \dfrac{\exp(x) - \exp(-x)}{\exp(x) + \exp(-x)}$.

For this case, one can show formally that if $J < H$, then the equilibria is unique whereas if $J > H$ there are three equilibria, of which only the two extremal equilibria (in terms of the magnitude of m_g) are stable under dynamic analogs of the model.

Comments

1. There is an interplay of the degree of unobserved heterogeneity and the strength of social interactions that determines the number of equilibria.

2. This is an example of a phase transition

3. The threshold for multiplicity depends on the number of choices.

Multinomial Choice Under Alternative Error Assumptions

The basic logic of the multinomial model is straightforward to generalize. This can be seen if one considers the preference structure

$$V_{i,l} = h_{i,l} + Jp_{i,l}^e + \beta^{-1}\varepsilon_{i,l} \qquad (24)$$

This is the same preference structure we worked with earlier, except that β is now explicitly used to index the intensity of choice (in the McFadden sense) rather than as a parameter of the distribution of the random payoff term $\varepsilon_{i,l}$.

We assume that these unobserved utility terms are independent and identically distributed with a common distribution function $F_{\varepsilon}(\cdot)$.

For this model, the probability that agent i makes choice l is

$$\mu\begin{pmatrix} \varepsilon_{i,0} - \varepsilon_{i,l} \le \beta\left(h_{i,l} - h_{i,0}\right) + \beta J\left(p_{i,l}^{e} - p_{i,0}^{e}\right),...,\\ \varepsilon_{i,L-1} - \varepsilon_{i,l} \le \beta\left(h_{i,l} - h_{i,L-1}\right) + \beta J\left(p_{i,l}^{e} - p_{i,L-1}^{e}\right) \end{pmatrix} \quad (25)$$

Following Anderson, dePalma, and Thisse (1992, pg. 36), conditional on a realization of $\varepsilon_{i,l}$, the probability that l is chosen is

$$\prod_{j \neq i} F_\varepsilon \left(\beta h_{i,l} - \beta h_{i,j} + \beta J p_{i,l}^e - \beta J p_{i,j}^e + \varepsilon_{i,l} \right) \qquad (26)$$

which immediately implies that the probability of the choice l without conditioning on the realization of $\varepsilon_{i,l}$ is

$$p_{i,l} = \int \prod_{j \neq l} F_\varepsilon \left(\beta h_{i,l} - \beta h_{i,j} + \beta J p_{i,l}^e - \beta J p_{i,j}^e + \varepsilon \right) dF_\varepsilon . (27)$$

Eqs. (16)-(19) provide a multinomial choice model whose structure is fully analogous to the multinomial logit structure developed under parametric assumptions. Under self-consistency, the aggregate choice probabilities of this general multinomial choice model are the solutions to

$$p_l = \int \int \prod_{j \neq l} F_\varepsilon \left(\beta h_l - \beta h_j + \beta J p_l - \beta J p_j + \varepsilon \right) dF_\varepsilon dF_h \quad (28)$$

As in the multinomial logit case, the compound parameter βJ plays a critical role in determining the number of self-consistent equilibrium choice probabilities p_l. This finding is formalized in Theorem 2.

Theorem 2. Uniqueness versus multiplicity of self-consistent equilibria in multinomial choice models with social interactions

Suppose that individual choices and associated self-consistent equilibria are described by (19)-(20). Assume that $h_{i,l} = 0$ $\forall i, l$ and $\varepsilon_{i,l}$ are independent across i and l. There exists a threshold T such that if $\beta J < T$, then there is a unique self-consistent equilibrium, whereas if $\beta J > T$ there exist at least three self-consistent equilibria.

The relationship between βJ and the number of equilibria is less precise than was found in Theorem 1, the multinomial logit case, as Theorem 3 does not specify anything about the way in which L, the number of available choices, affects the number of equilibria. This lack of precision is to be expected since we did not specify the distribution of the errors.

Groups Choice and Behavior Choice

Our analysis so far has treated groups as predetermined. For contexts such as ethnicity or gender this is presumably appropriate.

However, in other contexts, such as residential neighborhoods, group memberships are themselves presumably influenced by the presence of social interactions effects. Hence a complete model of the role of social interactions on individual and group outcomes requires a joint description of both the process by which neighborhoods are formed and the subsequent behaviors they induce.

A Nested Choice Approach to Integration of Behaviors and Group Memberships

A second approach to endogenizing group memberships may be developed using the nested logit framework originated by Ben Akiva (1973) and McFadden (1978). The basic idea of this framework is the following. An individual is assumed to make a joint decision of a group $g \in \{0,...G-1\}$ and a behavior $l \in \{0,...L-1\}$. We will denote the group choice of i as δ_i.

The structure of this joint decision is nested in the sense that the choices are assumed to have a structure that allows one to decompose the decisions as occurring in two stages: first, the group is chosen and then the behavior.

The key feature of this type of model is the assumption that choices at each stage obey a multinomial logit probability structure. For the behavioral choice, this means that

$$\mu\left(\omega_i = I \middle| h_{i,l,g}, p_{i,l,g}^e, \delta_i = g\right) = \frac{\exp \beta\left(h_{i,l,g} + Jp_{i,l,g}^e\right)}{\sum_{j=0}^{L-1} \exp \beta\left(h_{i,l,g} + Jp_{i,l,g}^e\right)} \quad (29)$$

which is the same behavioral specification as before.

Group membership choices are somewhat more complicated. In the nested logit model, group choices are assumed to obey

$$\mu\left(i \in g \middle| h_{i,l,g}, p_{i,l,g}^{e} \forall l, g\right) = \frac{\exp(\beta_{g} Z_{i,g})}{\sum_{g} \exp(\beta_{g} Z_{i,g})} \qquad (30)$$

where

$$Z_{i,g} = E(\max_{l} h_{i,l,g} + J p_{i,l,g}^{e} + \varepsilon_{i,l,g}) \qquad (31)$$

A standard result (e.g. Anderson, de Palma and Thisse (1992, pg. 46)) is that

$$E\left(\max\left(h_{i,l,g} + Jp_{i,l}^e + \varepsilon_{i,l,g}\Big| h_{i,l,g}, p_{i,l,g}^e \forall l, g\right)\right) =$$
$$\beta^{-1}\log\left(\sum_l \exp \beta\left(h_{i,l,g} + Jp_{i,l,g}^e\right)\right) \qquad (32)$$

Combining equations, the joint group membership and behavior probabilities for an individual are thus described by

$$\mu\left(\omega_i = l, \delta_i = n \middle| h_{i,l,n}, p_{i,l,n}^e \,\forall l, n\right) =$$

$$\frac{\exp\left(\beta_n\beta^{-1}\log\left(\sum_l \exp\beta\left(h_{i,l,n} + Jp_{i,l,n}^e\right)\right)\right)}{\sum_n \exp\left(\beta_n\beta^{-1}\log\left(\sum_l \exp\beta\left(h_{i,l,n} + Jp_{i,l,n}^e\right)\right)\right)} \cdot$$

$$\frac{\exp\beta\left(h_{i,l,n} + Jp_{i,l,n}^e\right)}{\sum_{j=0}^{L-1} \exp\beta\left(h_{i,l,n} + Jp_{i,l,n}^e\right)} \qquad (33)$$

This probabilistic description may be faulted in that it is not directly derived from a utility maximization problem. In fact, a number of papers have identified conditions under which the probability structure is consistent with utility maximization, cf. McFadden (1978) and Borsch-Supan (1990) for discussion. A simple condition (cf. Anderson, dePalma, and Thisse, 1992, pg. 48) that renders the model compatible with a well posed utility maximization problem is $\beta_n \leq \beta$, which in essence requires that the dispersion of random payoff terms across groups is lower than the dispersion in random payoff terms across behavioral choices within a group.

There has yet to be any analysis of models such as (36) when self-consistency is imposed on the expected group choice percentages $p_{i,l,g}^{e}$. Such an analysis should provide a number of interesting results. For example, a nested structure of this type introduces a new mechanism by which multiple equilibria may emerge, namely the influence of beliefs about group behaviors on group memberships, which reciprocally will affect behaviors. This additional channel for social interactions, in turn, raises new identification questions.

Comment: Existence may require membership "prices"

State of Literature

1. Social interactions may be integrated into standard choice models in ways that preserve neoclassical reasoning, yet allow for phenomena such as multiple equilibria

2. Much left to do, especially for nested choice generalizations that integrated group formation and behavior in groups. Social networks are even harder to integrate.