Social interactions in small groups

Brian Krauth
Simon Fraser University
Burnaby BC V5A 1S6
Email: bkrauth@sfu.ca
Web: http://www.sfu.ca/~bkrauth

August 26, 2003

Abstract

Most models of endogenous social interactions implicitly assume that individuals are interacting directly or indirectly with a very large (infinite) group of others. These models generally feature multiple equilibria when the marginal impact of peers exceeds some strictly positive critical value. This paper shows that this result depends critically on the assumption that the size of the peer group is large, and analyzes the properties of an alternative model in which the peer groups are small (finite). I find that when peer groups are small, multiple equilibria exist for any positive degree of peer influence. A brief application to the variation in youth smoking rates by ethnicity demonstrates the potential interaction of group size and strength of peer effect in generating a wide range of equilibrium behavior.

JEL codes: Z13, D00
Keywords: Social interactions, neighborhood effects, peer effects

1 Introduction

This paper analyzes the equilibrium behavior of a large population of heterogeneous agents organized into small and non-overlapping peer groups. Each agent makes a binary choice subject to a strategic complementarity with his or her peers. The complementarity may be interpreted as a simple taste for conformity, a technological benefit to coordination, or the result of a social learning process. The model in this paper deviates from a now-standard model (Schelling 1978, Brock and Durlauf 2001) in the "social interactions" literature only in that the size of the effective peer group remains finite, even as the size of the population increases. Most previous work in the literature implicitly assumes that the effective peer group is very large, and uses limit approximations to characterize behavior. While there are many applications in which the influential peer group is clearly large, for example, in studies of neighborhood or school influences on behavior, there are also many applications in which the relevant peer group is small. For example, empirical researchers have investigated the influence of university roommates (Sacerdote 2001, Kremer and Levy 2001), siblings and close friends (Duncan, Boisjoly and Harris 2000), and classmates (Hoxby 2000) on various choices and outcomes. This paper finds that equilibrium behavior exhibits substantially different properties depending on peer group size.

The primary impact of group size is on uniqueness versus multiplicity of equilibria. With large peer groups, the equilibrium is unique provided that the strength of peer influence is below some strictly positive critical value. In contrast, with small[1] peer groups, there are multiple equilibria provided that peer influence is greater than zero. The range of average behavior which is consistent with equilibrium varies with both the peer group size and the strength of the peer effect.

This result is of potential interest because multiplicity of equilibria due to social interactions has been used to rationalize a wide variety of otherwise puzzling phenomena, in particular high degrees of variation in behavior across social groups and/or time which are difficult to explain by variation in fundamentals[2] Multiplicity of equilibria has also been used to argue for an equally wide variety of policy measures[3] aimed at pushing particular social groups into "better equilibria". While such interpretations of models with multiple equilibria are controversial, the presence of multiple equilibria in a simple static model is generally associated with interesting aggregate behavior – history dependence, high sensitivity to small changes in fundamentals, rapid changes in behavior across time

[1]Formally, I use "small" to mean that the number of agents in the peer group is finite and "large" to mean that the number is approaching infinity. Informally, the results can be interpreted as applying to the common meanings of the term - the behavior of the small-group model converges to the behavior of the large-group model as the group size increases.

[2]To provide a few examples, Glaeser and Scheinkman (2000) suggest that this type of model can be used to rationalize stock market crashes, the prevalence of Christianity in France and Buddhism in Thailand, the Great Depression, the variation in crime rates across cities, and the brief but overwhelming popularity of the Hula Hoop. Schelling (1978) suggests it can explain attendance at seminars, applause, double parking, fashion, and smoking. Granovetter (1978) discusses applications to voting behavior, strikes and riots.

[3]For example, Moffitt (2001) discusses school busing, housing mobility and scattered-site housing programs, mass-media antismoking campaigns, and welfare reform

and space – in richer dynamic versions of the model. As a result, establishing multiplicity in the static model can be a starting point for investigating more complex and realistic settings.

In general those researchers who have used multiple equilibria to rationalize observed behavior have paid little attention to verifying whether the necessary conditions for multiple equilibria are satisfied. By establishing additional conditions which affect whether social interactions induce multiple equilibria, this paper makes progress toward the eventual goal of placing greater empirical discipline on speculations about their applicability to rationalizing particular variations in group behavior. Combining the results in this paper with the previous results in the literature, it is clear that multiple equilibria require either strong peer effects or small peer groups. To demonstrate how the results can be applied to place discipline on theorizing, the paper includes a simple example: interpreting the high variation in youth smoking rates by race and across time in the United States. The application suggests that a wide range of equilibrium behavior is consistent with current estimates of peer effects in youth smoking if and only if the relevant peer group is quite small.

1.1 Related literature

The large-groups model used as a benchmark here is also sometimes called the "critical mass", "threshold", or "epidemic" model and has a long history in both economics and sociology. Influential early treatments include Schelling (1978), Granovetter (1978), and Crane (1991), and a recent article by Brock and Durlauf (2001) provides a somewhat more rigorous analysis. In both the benchmark model and the model developed here, agents are exogenously assigned into peer groups and each individual must make a binary choice. The incremental utility associated with each option is increasing in the fraction of peers that choose that option. For example, a young person may experience social pressure to smoke, drink, or engage in other risky behavior if his or her peers do so. Alternatively, there may be a stigma to dropping out of the labor market, welfare receipt, or tax evasion, but the stigma fades as more community members engage in that behavior. Critical mass models are thus similar to N-player coordination games, and both fit into the category of supermodular games. Despite this similarity, the critical mass model and other models in the literature on social interactions are distinct from coordination games in that agents are explicitly heterogeneous. The focus in this literature is on neither the strategic aspects of interaction nor the beliefs of players about the strategies chosen by others but rather on the relationship between the distribution of preferences in a large population and the equilibrium distribution of behavior in that population. In the absence of heterogeneity, the critical mass model reduces to either a standard coordination game or a trivial game in which all agents have the same strictly dominant strategy.

While interactions between peers are necessarily complex and dynamic, this paper follows much of the literature on the critical mass model by characterizing the Nash equilibria of a static game. While an imperfect description of reality, the Nash equilibria of the static game are closely related to the long-run behavior in richer dynamic settings. For example, Blume and Durlauf (forthcoming) analyze Brock and Durlauf's (2001) version of the critical mass model under standard evolutionary dynamics. They find that if the static version of the model has multiple equilibria, the sample path of the associated dynamic model spends long periods in a neighborhood of either of the two stable Nash equilibria, and asymptotically converges to whichever stable Nash equilibria is risk dominant. Furthermore, the relationship between the mean of the preference distribution and the mean choice level in the risk dominant equilibrium is highly nonlinear. Two groups with apparently similar characteristics can exhibit quite different long-run behavior. These results formalize Schelling's (1978) insight that the multiple Nash equilibria of the simple critical mass model can be interpreted in a variety of related ways – self-fulfilling expectations, self-enforcing conventions or social norms, "tipping" or rapid unexplained changes in average behavior, or discontinuous relationships between incentives and average choices – depending on the details of one's assumptions about the information held by agents and the exact dynamics of interaction. As a result, this paper focuses on multiple equilibria in the static model as a starting point, with dynamic analysis left for future research.

2 Model and results

Section 2.1 outlines the basic behavioral framework of the two models, using the notation and terminology of Brock and Durlauf (2001). Sections 2.2 and 2.3 consider two different models that share this behavioral framework. Section 2.2 reviews the standard model in which the peer group is large. Section 2.3 introduces the small group version of the model and derives its properties.

2.1 Choices and preferences

This section establishes the behavioral framework common to both the large group and small group models. A peer group consists of n players, indexed by $i \in \{1, 2, \dots, n\}$. Each player chooses an action ω_i from the pure strategy space $\{-1, 1\}$. Let $\omega \equiv (\omega_1, \omega_2, \dots, \omega_n)$ be a pure strategy profile for the peer group and let ω_{-i} be a strategy profile which excludes player i's choice. I also define a mixed strategy as $\sigma_i \in [0, 1]$ where $\sigma_i = \Pr(\omega_i = 1)$, and define σ and σ_{-i} as the mixed strategy equivalents of ω and ω_{-i}.

The players interact through an exogenous directed graph G, known as the social network. Formally, G is a collection of sets $\{G_1, G_2, \dots, G_n\}$ where G_i is the set of all players whose actions directly influence player i. This formulation of the social network includes both global interactions (in which every player is influenced by every other player) and local interactions (in which players are arranged in space and each player is influenced by all players within a particular distance) as special cases. Player i's payoff (expected utility) function is:

$$u_i(\omega_i \ ; \ \omega_{-i}) = (h + J\bar{m}_i + \epsilon_i)\, \omega_i \tag{A1}$$

where:

$$\bar{m}_i \equiv \begin{cases} \frac{1}{\#G_i} \sum_{j \in G_i} \omega_j & \text{if } G_i \neq \emptyset \\ 0 & \text{if } G_i = \emptyset \end{cases}$$

The components of this payoff function are as follows. The choice variable ω_i enters into the payoff function multiplicatively, so that player i will prefer $\omega_i = +1$ when $(h + J\bar{m}_i + \epsilon_i)$ is positive and $\omega_i = -1$ when it is negative. This particular functional form is chosen to facilitate comparison with Brock and Durlauf (2001).

The parameter h and exogenous variable ϵ_i describe the attractiveness to a player i of choosing $+1$ relative to that of choosing -1, taking the behavior of others as given. The ϵ_i term is an exogenous and observable random variable with an as-yet-unspecified probability distribution. A preference profile for the group is given by $\epsilon \equiv (\epsilon_1, \epsilon_2, \ldots, \epsilon_n)$. The purpose of ϵ is not to model incomplete information but rather to model heterogeneity in preferences or opportunities. As mentioned previously, the use of random utility terms to model heterogeneity is a distinguishing characteristic of the social interactions literature.

The $J\bar{m}_i$ term describes the social influence on player i's choice. The parameter J measures the strength of peer influence, and it is assumed that

$$J \geq 0 \tag{A2}$$

The endogenous variable $\bar{m}_i$ is simply the average choice[4] among some subset of player i's peers. By assumption (A2), the relative payoff to each option is nondecreasing in the fraction of peers who choose that option.

The game itself is defined in normal form, with players $i \in \{1, 2, \ldots, n\}$, pure strategy space $\{-1, +1\}^n$, and payoff functions given by (A1) and (A2).

2.2 Model #1: Large group interactions

This section reviews a version of the critical mass model developed by Brock and Durlauf (2001), and discusses a few of their results. Although the notation and functional forms are specific to their paper, the substance of the results presented here applies to the other critical mass models in the literature.

For analytical convenience, Brock and Durlauf assume that ϵ_i has an IID logistic distribution:

$$\Pr(\epsilon_i \leq x | \epsilon_{-i}) = \frac{1}{1 + e^{-2\beta x}} \tag{A3}$$

The parameter $\beta > 0$ indexes the the degree of homogeneity in preferences across group members. Homogeneity is increasing in β, i.e., as β approaches infinity the variance of ϵ_i approaches zero.

[4] This type of peer effect is often called an "endogenous" social interaction effect (Manski 2000) to distinguish it from a "contextual" effect, in which each agent is affected by the exogenous characteristics (for example socioeconomic status) of his or her peers.

They also assume that interactions are global, i.e., that each individual is influenced by the average choice in the peer group as a whole:

$$G_i = \{1, 2, \ldots, i-1, i+1, \ldots, n\} \tag{A4}$$

which implies that $\bar{m}_i = \frac{1}{n-1} \sum_{j \neq i} \omega_j$.

Rather than using a standard equilibrium concept, Brock and Durlauf characterize what they call a "self-consistent equilibrium," which is defined as a joint probability distribution over choices such that, for all i:

$$\omega_i = \begin{cases} +1 & \text{if } h + JE(\bar{m}_i) + \epsilon_i > 0 \\ -1 & \text{if } h + JE(\bar{m}_i) + \epsilon_i < 0 \end{cases} \tag{1}$$

Note that this condition assumes that players choose a best response to the unconditional expected value of average behavior in the group rather than the realized average behavior. The equilibrium condition (1) can be interpreted[5] as an approximate description of the Nash equilibria as n approaches infinity. As the size of the peer group approaches infinity, the average choice $\bar{m}_i$ converges in probability to its expected value $E(\bar{m}_i)$. It is in this sense that the Brock-Durlauf model and other critical mass models[6] are "large group" models.

Given the model and definition of equilibrium, Brock and Durlauf demonstrate several results. First, self-consistent equilibrium implies that:

$$E(\omega_i) = \tanh(\beta h + \beta JE(\bar{m}_i)) \tag{2}$$

The tanh function, not well known in economics, appears frequently in the sta-

[5] There is an alternative interpretation of equation (1) as describing the Nash equilibria of an n-player one-shot simultaneous move game in which agents do not observe one another's type (value of ϵ_i). In the absence of such information, agents choose a best response to the expected value of $\bar{m}_i$ in the population. Other work by both authors (Brock 1993, Durlauf 1997, Blume and Durlauf forthcoming) as well as work by other authors extending the model (Ioannides 2001, Glaeser and Scheinkman 2000), generally uses the large groups interpretation rather than the uninformed agents interpretation. This is most likely because a peer group almost by definition consists of individuals engaging in repeated interaction and exchange of information, so applications with uninformed agents and one-shot choices are less common than applications with large groups. This paper also follows the large groups interpretation of the standard critical mass model.

[6] Other discussions are even less explicit in the assumption of a large group. For example, Schelling (1978, p. 102) states "We idealize the bar chart [empirical distribution of preferences] into a smooth frequency distribution" without noting that this approximation implicitly assumes large groups.

tistical mechanics literature from which this model is adapted. It is continuous, differentiable, and strictly increasing, and takes on values in the interval $(-1, +1)$. Differentiating the right side of equation (2) gives the marginal impact of average peer choice on an individual's choice:

$$\frac{\partial E(\omega_i)}{\partial E(\bar{m}_i)} = \beta J \left(\text{ sech } (\beta h + \beta J E(\bar{m}_i))\right)^2 \leq \beta J \tag{3}$$

Because equation (2) holds for every i, rational expectations implies that $E(\omega_i) = E(\omega_j) = E(\bar{m}_i) = E(\bar{m}_j) = m^*$, where m^* is a solution to the equation:

$$m^* = \tanh(\beta h + \beta J m^*) \tag{4}$$

m^* can be interpreted as the expected average choice in the group and $\frac{m^*+1}{2}$ as the expected proportion choosing $\omega_i = 1$. By the previously described properties of the tanh function, a solution to equation (4) always exists. A sufficient condition for uniqueness is that:

$$\frac{\partial E(\omega_i)}{\partial E(\bar{m}_i)} < 1 \text{ for all } E(\bar{m}_i) \in [-1, 1] \tag{5}$$

Glaeser and Scheinkman (2000) refer to (5) as the "Moderate Social Influence" (MSI) condition and note that it implies uniqueness in a wide class of social interaction models. For the specific model presented here, equation (3) implies that the MSI condition holds if $\beta J < 1$. When the MSI condition does not hold, there can be multiple equilibria:

Proposition 1 (Brock-Durlauf (2001) Proposition 2) *1. If $\beta J > 1$ and $h = 0$, there exist three roots to equation (4). One of these roots is positive, one root is zero, and one root is negative.*

2. If $\beta J > 1$ and $h \neq 0$, there exists a threshold H (which depends on βJ) such that

   (a) For $|\beta h| < H$, there exist three roots to equation (4), one of which has the same sign as h, and the others possessing opposite sign.

   (b) For $|\beta h| > H$, there exists a unique root to equation (4), with the same sign as h.

When equilibrium is not unique, the average behavior of a social group is not uniquely determined by economic fundamentals. Under the assumptions of the model, Proposition 1 gives conditions which are associated with multiple equilibria: influential peers (βJ large) and small differences between choices (βh close to zero). The following two sections identify a third factor: small peer groups.

2.3 Model #2: Small-group interactions

The small-group model is nearly identical to the standard critical mass model described in the previous section. However, unlike the standard model's implicit assumption that the peer group is large, the small-group model has explicitly finite n and uses standard equilibrium concepts.

First note that the game falls within the class of supermodular games (Milgrom and Roberts 1990), the properties of which are both convenient and well understood. In particular, there exist both sharp and easily calculated upper and lower bounds on equilibrium behavior that apply regardless of the equilibrium concept used. These properties are summarized in Proposition 2.

Proposition 2 (Properties of equilibria) *Consider an n-player game in normal form with pure strategy space $\{-1, +1\}$ and payoff functions described by (A1) and (A2). For a given ϵ, there exist pure strategy profiles $\omega^L(\epsilon)$ and $\omega^H(\epsilon)$ such that:*

1. *All serially undominated strategy profiles lie in the interval[7] $[\omega^L(\epsilon), \omega^H(\epsilon)]$. This also implies that all rationalizable strategy profiles, all correlated equilibria and all Nash equilibria lie in this interval.*
2. *Both $\omega^L(\epsilon)$ and $\omega^H(\epsilon)$ are Nash equilibria. This also implies that they are correlated equilibria, rationalizable, and serially undominated.*
3. *Both $\omega^L(\epsilon)$ and $\omega^H(\epsilon)$ are nondecreasing in h.*

Milgrom and Roberts also show that a wide class of learning dynamics (including best-response and fictitious play) converge to behavior in the $[\omega^L(\epsilon), \omega^H(\epsilon)]$ interval. These results taken together are particularly powerful because they imply

[7] Formally, a pure strategy profile ω lies in the interval $[x, y]$ if and only if $x_i \leq \omega_i \leq y_i$ for all i. A mixed strategy profile lies in $[x, y]$ if and only if any pure strategy profile in its support lies in $[x, y]$.

that one can easily find (through iterated removal of strictly dominated strategies) sharp upper and lower bounds on behavior that do not depend on one's preferred solution concept.

Next, these upper and lower bounds on behavior for a particular group with a particular preference profile ϵ are used to construct upper and lower bounds on aggregate behavior as a function of the model parameters. In doing so, I follow Jovanovic (1989) in interpreting an economic model as imposing a set of restrictions on the conditional probability distribution of endogenous variables (ω) given model parameters (h, J) and exogenous variables (ϵ). These restrictions can then be integrated over the probability distribution of exogenous variables to generate aggregate predictions.

In particular, we wish to characterize the restrictions the model places on the average choice in a large population. To facilitate this, an equilibrium average choice is defined as a quantity $m^* \in [\bar{m}^L, \bar{m}^H]$ where:

$$\begin{aligned} \bar{m}^L &\equiv E\left(\frac{1}{n}\sum_{i=1}^{n}\omega^L(\epsilon)\right) \\ \bar{m}^H &\equiv E\left(\frac{1}{n}\sum_{i=1}^{n}\omega^H(\epsilon)\right) \end{aligned} \tag{6}$$

If $\bar{m}^L = \bar{m}^H$, the equilibrium average choice is unique. If the equilibrium average choice is not unique ($\bar{m}^L < \bar{m}^H$), then the model is consistent with a variety of predictions for the average prevalence of a given behavior in a population. Note that the existence of multiple equilibria for some values of ϵ is necessary but not sufficient for the existence of multiple values for equilibrium average choice. In addition, it is necessary for those values of ϵ that imply multiple equilibria to occur with strictly positive probability.

We look at two versions of the model. The first version is quite general. Let G be an arbitrary directed graph, and let the probability distribution of ϵ be an arbitrary distribution such that:

$$\forall x \in R^n, \delta > 0, \qquad \begin{array}{c} \Pr(\epsilon = x) = 0 \\ \Pr(||x - \epsilon|| < \delta) > 0 \end{array} \tag{A3'}$$

It is also useful to define the graph-theoretic notion of a cycle in the context of this model. A cycle is a sequence of players C of length $K > 1$ such that $C_1 = C_K$ and $C_{k+1} \in G_{C_k}$ for all $k < K$. In other words, a cycle is a chain of social influences that return to their starting player.

Proposition 3 (Multiple equilibria) *Consider an n-player game in normal form with pure strategy space $\{-1, +1\}$ and payoff functions described by (A1), (A2), and (A3'). The equilibrium average choice is unique if $J = 0$ or if G has no cycles. It is not unique if $J > 0$ and G has at least one cycle.*

The second version of the model is a small-groups variation on Brock and Durlauf's model, which maintains the assumption of global interactions within the peer group and an IID logistic distribution for preferences, but replaces the implicit assumption of a large peer group with an explicitly finite peer group.

Proposition 4 (Multiple equilibria w/Brock-Durlauf assumptions) *Consider an n-player game in normal form with pure strategy space $\{-1, +1\}$ and payoff functions described by (A1)-(A4). The equilibrium average choice is unique if $\beta J = 0$ and is not unique if $\beta J > 0$.*

Propositions 3 and 4 say that in a large population of small peer groups, the existence of any endogenous peer effect, regardless of its strength, is sufficient for average behavior in the population to not be uniquely determined by preferences. This is in contrast to the results for the large groups model, which tend to imply a strictly positive critical value for the peer effect below which average behavior is uniquely determined by preferences.

The intuition for these results is simple. First, the existence of a cycle in the network is necessary in order for the possibility of a self-justifying shift in behavior by a subset of players. For example, when there is a cycle of length two there will exist preference profiles such that it is a best reply for player one to choose whatever player two chooses, and vice versa. The existence of at least one cycle is a fairly weak requirement on a social network. For example, every social network in which each player has at least one influential peer (i.e., G_i is nonempty for all i) has a cycle. One example of a social network with no cycles is the star or leader-follower network, in which everyone is influenced only by the behavior of a leader and the leader is influenced by no one. Here, the equilibrium is unique (with probability one) regardless of the strength of peer effects - the leader has a strictly dominant strategy, and each follower has a dominant strategy after the dominated strategy for the leader has been eliminated.

Second, it should be noted that regardless of the group size, there exist at least some preference profiles that are consistent with multiple equilibria. For example, for ϵ in a neighborhood around $(-h, -h, \ldots, -h)$ we have $\omega^L(\epsilon) = (-1, -1, \ldots, -1)$ and $\omega^H(\epsilon) = (+1, +1, \ldots, +1)$, regardless of the value of n. However, as n approaches infinity, the probability that ϵ will be in that neighborhood approaches zero.

Whether the equilibria of the model differ in a quantitatively interesting manner depends on the parameter values. Section 3 discusses the quantitative properties of these models in the context of a simple application.

3 Application: The youth smoking gap

This section discusses one possible application of the model: interpreting the vast difference in smoking rates between black and white youth in the United States during the 1990's. Figure 1 shows that although both groups had similar smoking rates in the 1970's, by the early 1990's the smoking rate among black youth was one-third the smoking rate of white youth. This gap appears in multiple surveys during this period and is robust to corrections for possible selection or under-reporting bias, as well as to controlling for standard covariates such as prices, disposable income, state-level policies, family characteristics and behavior, etc. (Gruber and Zinman 2001). Durlauf and Young (2001, p. 6) suggest that the critical mass model can potentially explain the phenomena depicted in Figure 1, i.e., that successive cohorts of black youth switched from a high-smoking equilibrium to a low-smoking equilibrium, while white youth remained in the high-smoking equilibrium.

This interpretation raises the question of whether the necessary conditions for generating a wide range of equilibrium behavior are satisfied. This section uses a simple numerical experiment to provide a tentative answer. The methodology is

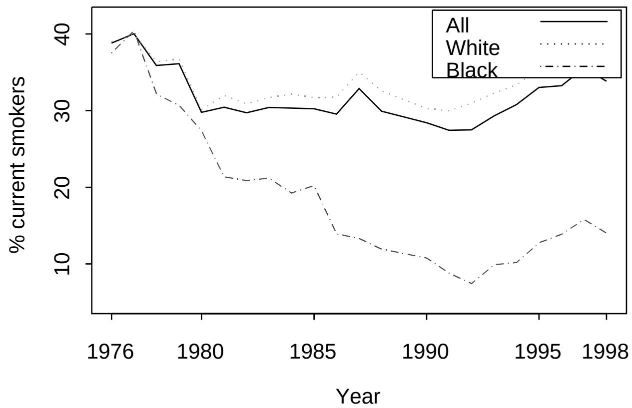

Figure 1: Self-reported current smoking rates for white and black youth in the United States, 1976-1998. Source: Monitoring the Future Survey (Johnston et al., 1998).

to assume that the world corresponds to the model described in assumptions (A1)-(A4), select parameter values based on previous empirical studies, solve the model numerically, and compare the range of equilibrium smoking rates implied by the model to the range observed across the two groups in the data. It should be emphasized that the primary purpose of this analysis is to illustrate a potential application of the model, and not to provide a definitive answer to the substantive question. A number of issues regarding the reliability of existing parameter estimates, the applicability of the static model developed here to a clearly dynamic process of smoking adoption, and the lack of a theory of equilibrium selection imply that a complete treatment of the issue must be left to future research.

The parameters to be selected are the marginal peer effect (J), the average benefit from smoking (h), and the size of the group (n). Only βJ and βh are identified in behavioral data, so I follow the standard practice from the discrete choice literature and normalize β to one. The group sizes selected are 2 (best-friend interactions), 6 (close friend interactions - results from multiple surveys indicate that both young people and adults usually report having 5 to 7 close friends), 30 (classmate interactions), and infinity (neighborhood interactions). The private incentive to smoke (h) is set so that the smoking rate implied by $\bar{m}^H$ is approximately 30%, i.e., the approximate smoking rate for white youth in the early 1990's. The marginal peer effect J is set on the basis of existing empirical studies. Most empirical studies of peer effects are flawed by their failure to account for the endogeneity of peer behavior (Manski 2000); however three recent papers use an instrumental variables approach to estimate peer effects in youth smoking while accounting for endogeneity. Each uses a slightly different model, each imposes strong and potentially untenable identifying assumptions, and none estimate a model that is directly applicable to the model developed here. As a result, their estimates are used only as a rough guide in picking a reasonable value to use for the marginal peer effect J. Gaviria and Raphael (2001, Table 5) find a marginal impact of peer smoking of 0.156, i.e., a one percentage point increase in smoking among a young person's peers is associated with a 0.156 percentage point increase in the probability that the young person will smoke. Norton, et al. (1998, Table 6) find a marginal impact of 0.942, and Powell, et al. (2003, Table 2) find a marginal impact of 0.5385. For the numerical experiments with the model, I use the intermediate value of $J = 0.5$, which implies (according to equation (3)) a marginal impact of up to 0.5.

Given the parameters, I solve for the range of equilibrium average choice $[\bar{m}^L, \bar{m}^H]$. For $n \to \infty$, the model can be solved directly by solving equation (4). For finite n, the endpoints of the range $[\bar{m}^L, \bar{m}^H]$ can be approximated using simulation. In each simulation run $s \in \{1, 2, \ldots, S\}$, the preference profile ϵ^s is drawn from the IID logistic distribution and the pure strategy equilibria $\omega^L(\epsilon^s)$ and $\omega^H(\epsilon^s)$ (as defined in Proposition 2) are found. Estimates of $\bar{m}^L$ and $\bar{m}^H$ are found by averaging across players and simulations:

$$
\begin{aligned}
\bar{m}^L &\approx \frac{1}{S}\sum_{s=1}^{S}\left(\frac{1}{n}\sum_{i=1}^{n}\omega_i^L(\epsilon^s)\right) \\
\bar{m}^H &\approx \frac{1}{S}\sum_{s=1}^{S}\left(\frac{1}{n}\sum_{i=1}^{n}\omega_i^H(\epsilon^s)\right)
\end{aligned} \tag{7}
$$

Table 1 shows the results of simulating the model 100,000 times for the selected parameter values and for different values of the group size n. This range can be compared to the range of actual behavior observed in Figure 1.

These results show how the range of equilibria depends on the size of the peer group as well as on the strength of the peer effect. When individuals are only influenced by their best friend ($n = 2$), equilibrium smoking rates range from 16% to 30%. As the size of the group increases the range of equilibrium smoking rates gradually narrows. Finally, when individuals are members of very large peer groups ($n \to \infty$), the equilibrium smoking rate is unique because $\beta J = 0.5 < 1$. We can compare these results to the actual gap in smoking rates by race of roughly twenty percentage points. If the model is correctly specified, a substantial fraction of the gap in behavior could be attributable to white and black youth playing different equilibria as long as the relevant peer group is small. If the group is sufficiently large there is no significant gap in behavior that could be attributed to multiple equilibria. This suggests, in keeping with the suggestions of Manski (2000), that empirical researchers should pay more serious attention to determining not just how much a young person is influenced by his or her peers, but also identifying what types of peers - best friends, close friends, classmates, schoolmates, etc. - are most influential.

Size of peer group (n)	Private incentive (h)	Estimated $[\bar{m}^L, \bar{m}^H]$	Range of equilibrium smoking rates
2	-0.514	$[-0.68, -0.4]$	$[16\%, 30\%]$
5	-0.335	$[-0.54, -0.4]$	$[23\%, 30\%]$
30	-0.250	$[-0.44, -0.4]$	$[28\%, 30\%]$
∞	-0.047	-0.4	30%

Table 1: Range of equilibrium smoking rates implied by the model.

4 Conclusion

The belief that models with multiple equilibria are a useful way of understanding large scale social change has often been controversial. Abstracting from the more fundamental theoretical criticisms of models with multiple equilibria (Morris and Shin 2001, for example), there is also an empirical question of whether the necessary conditions for the presence of multiple equilibria are satisfied.

The results in this paper make some progress toward answering this question in the context of binary choice. When the relevant peer group is large, so that each individual has a very small impact on the choices of others, equilibrium is unique for moderate degrees of social influence, but not unique if social influence is strong enough. When the relevant peer group is small, and when peers have mutual influence on one another in the sense that the social network has a cycle, the equilibrium average choice is not unique as long as the peer effect is positive. The range of equilibria narrows both as peer groups get larger and as the peer effect decreases.

Applied to specific cases, this suggests that multiple equilibria may be of less interest when considering interactions that occur at the level of the neighborhood or school, and more interest when considering interactions between roommates or close friends. One issue that has not yet received much attention in the empirical work is determining at what level of aggregation peer influence operates. The results here suggest this question is deserving of more attention in the future.

References

Blume, Lawrence and Steven N. Durlauf, "Equilibrium concepts for social interaction models," *International Game Theory Review*, forthcoming.

Brock, William A., "Pathways to randomness in the economy: Emergent nonlinearity and chaos in economics and finance," *Estudios Economicos*, 1993, *8*, 3–55. Also available as SSRI Reprint #410, University of Wisconsin.

— **and Steven N. Durlauf**, "Discrete Choice with Social Interactions," *Review of Economic Studies*, 2001, *68*, 235–260.

Crane, Jonathan, "The Epidemic Theory of Ghettos and Neighborhood Effects on Dropping Out and Teenage Childbearing," *American Journal of Sociology*, 1991, *96*, 1226–1259.

Duncan, Greg J., Johanne Boisjoly, and Kathleen Mullan Harris, "Sibling, peer, neighbor, and schoolmate correlations as indicators of the importance of context for adolescent development," *Demography*, 2000, *38*, 437–447.

Durlauf, Steven N., "Statistical mechanics approaches to socioeconomic behavior," in W. Brian Arthur, Steven N. Durlauf, and David A. Lane, eds., *The Economy as an Evolving Complex System II*, Addison-Wesley, 1997, pp. 81–104.

— **and H. Peyton Young**, "The new social economics," in Steven N. Durlauf and H. Peyton Young, eds., *Social Dynamics*, MIT Press, 2001, pp. 1–14.

Gaviria, Alejandro and Steven Raphael, "School-based peer effects and juvenile behavior," *Review of Economics and Statistics*, 2001, *83* (2), 257–268.

Glaeser, Edward L. and José A. Scheinkman, "Non-market interactions," Working Paper 8053, NBER 2000.

Granovetter, Mark, "Threshold models of collective behavior," *American Journal of Sociology*, 1978, *83* (6), 1420–1443.

Gruber, Jonathan and Jonathan Zinman, "Youth smoking in the U.S.: Evidence and implications," in Jonathan Gruber, ed., *Risky Behavior Among Youths: An Economic Analysis*, University of Chicago Press, 2001.

Hoxby, Caroline M., "Peer effects in the classroom: Learning from gender and race variation," Working Paper 7867, NBER 2000.

Ioannides, Yannis M., "Topologies of social interactions," Working Paper, Tufts University 2001.

Johnston, Lloyd D., Jerald G. Bachman, Patrick M. O'Malley, John E. Schulenberg, and John Wallace, *Monitoring the Future Study, 1975-1998: United States*, Ann Arbor, MI: Institute for Social Research, University of Michigan, 1998.

Jovanovic, Boyan, "Observable implications of models with multiple equilibria," *Econometrica*, 1989, *57* (6), 1431–1437.

Kremer, Michael and Dan M. Levy, "Peer effects in alcohol use among college students," Working Paper, Harvard University 2001.

Manski, Charles F., "Economic Analysis of Social Interactions," *Journal of Economic Perspectives*, 2000, *14*, 115–136.

Milgrom, Paul and John Roberts, "Rationalizability, learning, and equilibrium in games with strategic complementarities," *Econometrica*, 1990, *58*, 1255–1277.

Moffitt, Robert A., "Policy interventions, low-level equilibria, and social interactions," in Steven N. Durlauf and H. Peyton Young, eds., *Social Dynamics*, MIT Press, 2001, pp. 45–82.

Morris, Stephen and Hyun Song Shin, "Rethinking multiple equilibria in macroeconomic modelling," in Ben S. Bernanke and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2000*, MIT Press, 2001.

Norton, Edward C., Richard C. Lindrooth, and Susan T. Ennett, "Controlling for the endogeneity of peer substance use on adolescent alcohol and tobacco use," *Health Economics*, 1998, pp. 439–453.

Powell, Lisa M., John A. Tauras, and Hana Ross, "Peer effects, tobacco control policies, and youth smoking behavior," Working Paper, University of Illinois-Chicago 2003.

Sacerdote, Bruce, "Peer effects with random assignment: Results for Dartmouth roommates," *Quarterly Journal of Economics*, May 2001, *116* (2), 655–680.

Schelling, Thomas, *Micromotives and Macrobehavior*, W.W. Norton, 1978.

A Proofs of propositions

Proposition 2: Properties of equilibria

First we show that the game is supermodular by the definition of Milgrom and Roberts (1990). This requires that (1) the strategy space of each agent is a complete lattice, (2) $u_i(\omega_i, \omega_{-i})$ is order upper semi-continuous in ω_i, order upper continuous in ω_{-i}, and has a finite upper bound, (3) u_i is supermodular in ω_i, and (4) u_i exhibits increasing differences in both arguments. First note that $\Omega \equiv \{-1, +1\}$ is finite. Therefore it has a well-defined supremum and infimum, for any subset $T \subset \Omega$, we have $\sup(T), \inf(T) \in \Omega$, so Ω is a complete lattice. The finiteness of Ω also implies that $u_i(\omega_i, \omega_{-i})$ is continuous in both arguments and has a finite upper bound. For any ω_i, ω_i', we have $u_i(\omega_i, \omega_{-i}) + u_i(\omega_i', \omega_{-i}) = u_i(\sup(\omega_i, \omega_i'), \omega_{-i}) + u_i(\inf(\omega_i, \omega_i'), \omega_{-i})$, so u_i is supermodular in ω_i. Finally, note that $u_i(+1, \omega_{-i}) - u_i(-1, \omega_{-i}) = 2(h + J\bar{m}_i + \epsilon_i)$. Since $\bar{m}_i$ is increasing in ω_{-i} and $J \geq 0$, this implies that u_i exhibits increasing differences in both ω_{-i} and in h. The proposition then follows directly from Milgrom and Roberts (1990), Theorems 5 and 6.

Proposition 3: Multiple equilibria

The proof that equilibrium average choice is unique when $J = 0$ is trivial. If $h+\epsilon_i \neq 0$ then player i has a dominant strategy, and $\Pr(h+\epsilon_i \neq 0) = 1$. Therefore with probability one the game has a unique dominant strategy equilibrium, so equilibrium average choice is unique.

The proof for uniqueness when G contains no cycles is as follows. First, note that a standard result in graph theory is that every acyclic digraph has at least one node of outdegree zero, i.e., if G has no cycles then there is at least one i such that $G_i = \emptyset$. This result will be used to demonstrate that the game is dominance solvable with probability one. Let $N = \{1, 2, \ldots, n\}$ be the set of players. Let $Z_1 \equiv \{i : G_i = \emptyset\}$. With probability one, each player in Z_1 has a strictly dominant strategy. Eliminate all dominated strategies for those players. Now, for any integer $k > 1$ define $Z_k \equiv \{i : G_i \subset Z_{k-1}\}$. Z_k is the set of all players who are influenced only by players in set Z_{k-1}. Two things should be noted. First, if each player in Z_{k-1} has a unique strategy that survives iterative removal of strictly dominated strategies, then each player in Z_k also has a unique such strategy with probability one. Second, because any subgraph of G is also acyclic, Z_k will have at least one more member than Z_{k-1} unless $Z_{k-1} = N$. Applying induction, with probability one, every player in Z_k has only one strategy that satisfies iterated strict dominance, and $Z_N = N$. The game is therefore dominance solvable, with probability one, and thus has a unique equilibrium average choice.

The proof for nonuniqueness when $J > 0$ and G contains at least one cycle is as follows. Note that:

$$\bar{m}^H - \bar{m}^L = E\left(\frac{1}{n}\sum_{i=1}^{n} \omega^H(\epsilon) - \omega^L(\epsilon)\right) \tag{8}$$

As $\omega^H(\epsilon) \geq \omega^L(\epsilon)$ for all ϵ it is sufficient to show that there exists a positive probability set of values for ϵ that imply $\omega_i^H(\epsilon) > \omega_i^L(\epsilon)$ for at least one agent i. Pick an agent j who is a member of one of the cycles and then pick the shortest cycle C of which j is a member. If there are ties, pick any cycle among the shortest. Let k be the length of C. Let

$$A = \left\{\epsilon : \begin{array}{ll} \epsilon_i \in (1-h, \infty) & \text{if } i \notin C \\ \epsilon_i \in (0, 1/n) & \text{if } i \in C \end{array}\right\}$$

Again, the assumptions made on the probability distribution of ϵ imply that $\Pr(\epsilon \in A) > 0$. By inspection, for all agents not in C, choosing -1 is a strictly dominant strategy. Also by inspection it is a best response for all agents in C to choose -1 if all other agents in C do so, and it is a best response for all agents in C to choose $+1$ if all other agents in C do so. Therefore $\bar{m}^H > \bar{m}^L$.

Proposition 4: Multiple equilibria with Brock-Durlauf assumptions

First, note that the graph described in (A4) has many cycles. To pick one, note that $1 \in G_2$ and $2 \in G_1$, so $(1, 2, 1)$ is a cycle. Second, note that the distribution described in (A3) satisfies condition (A3'). The result then follows directly from Proposition 3.