Multistable binary decision making on networks

Andrew Lucas*
Jefferson Physical Laboratory, Harvard University, Cambridge, Massachusetts 02138, USA and Department of Physics, Stanford University, Stanford, California 94305, USA

Ching Hua Lee
Department of Physics, Stanford University, Stanford, California 94305, USA
(Received 9 December 2012; published 12 March 2013)

We propose a simple model for a binary decision making process on a graph, motivated by modeling social decision making with cooperative individuals. The model is similar to a random field Ising model or fiber bundle model, but with key differences in behavior on heterogeneous networks. For many types of disorder and interactions between the nodes, we predict with mean field theory discontinuous phase transitions that are largely independent of network structure. We show how these phase transitions can also be understood by studying microscopic avalanches and describe how network structure enhances fluctuations in the distribution of avalanches. We suggest theoretically the existence of a "glassy" spectrum of equilibria associated with a typical phase, even on infinite graphs, so long as the first moment of the degree distribution is finite. This behavior implies that the model is robust against noise below a certain scale and also that phase transitions can switch from discontinuous to continuous on networks with too few edges. Numerical simulations suggest that our theory is accurate.

DOI: 10.1103/PhysRevE.87.032806 PACS number(s): 89.75.Hc, 46.65.+g, 89.65.−s

I. INTRODUCTION

Over the past decade there has been an explosion of interest in the statistical physics community in the behavior of simple statistical models on networks. These models often lead to insight about qualitative behavior in social systems as random graphs are a low-order approximation to realistic social networks [1,2]. Decision making processes have long been studied as a simple example of such an application. The voter model [3], along with variations with nonlinearities [4] or other complications [5,6], is a famous example, although it is only a model of consensus building. The Axelrod model is an alternative model that exhibits equilibria with a diversity of opinions [7–10]. Other spin models or agent-based models have been proposed to study financial markets [11,12].

Many of the above models do not predict a key phenomenon: the presence of shocks, catastrophes, and discontinuous phase transitions as external parameters are slowly tuned. Oftentimes, entirely new models have been proposed to account for this phenomenon [13–15]. More interesting, however, is the proposal that the random field Ising model, well known for hysteresis and discontinuous phase transitions [16], can be used to model this phenomenon in social science [17–19]. A similar model called the fiber bundle model, used to study the breakdown of some materials, also has similarly promising features [20,21]. Similarities between the fiber bundle model at a phase transition and the behavior of financial markets have also been noted [22]. Other recent models have attempted to discuss disorder-induced phase transitions of opinion dynamics, using disorder in the interactions between individuals [23].

In this paper we propose a very simple model for a binary decision making process on a network. Our model is similar to the random field Ising model in a global magnetic field, but with some important differences that make our model nearly exactly solvable on heterogeneous graphs. A preliminary mean field analysis of the model predicts disorder-induced discontinuous phase transitions and hysteresis. We then provide a microscopic justification for mean field theory as well, describing the microscopic dynamics of binary decision making in terms of avalanches. We also describe how fluctuations in the sizes of avalanches can scale with the size of the network on certain heterogeneous networks with fat tails.

Most interestingly, we will show that there is an infinite spectrum of equilibria in the large graph limit. This is not surprising because the random field Ising model has spin glasslike characteristics, often enhanced on networks [24–26]. Since our model does not admit a Hamiltonian and free energy, the glassy behavior of the binary decision model will be characterized by the presence of this spectrum of equilibria. We will then use this spectrum of equilibria to justify two phenomena: the robustness of the binary decision model to small fluctuations and the possibility that network structure can suppress a discontinuity in the phase transition.

The outline of our paper is as follows. In Sec. II we describe the binary decision model and provide some intuitive justification. In Sec. III we describe a mean field analysis of the model, beginning with an exactly solvable case that has discontinuous phase transitions. We then discuss numerical simulations, confirm that the model is roughly independent of the network structure, and discuss fluctuations in equilibria due to small network sizes. Section IV describes avalanche dynamics and provides a microscopic explanation for the independence of the mean field theory on network structure. In Sec. V we describe multistability, which is the appearance of an exponential number of equilibria. We propose this phenomenon first through a heuristic argument and then through a more rigorous cavity calculation. We conclude the paper with a discussion of basic consequences of multistability.

*lucas@fas.harvard.edu

II. BINARY DECISION MODEL

In this section we introduce the binary decision model, justifying its use as a simple model for equilibrium social behavior. We begin by approximating that individuals interact via a social network, which can be described as an undirected graph G consisting of a vertex set V and edges E. We define $N = |V|$, i.e., there are N nodes in the graph. The number of edges of a given node v in the graph is denoted by k_v; in mean field theory, we will often group together all nodes with the same number of edges, as is often done [1]. With each node v in the graph, we associate a binary variable $x_v \in \{0,1\}$. For example, $x_v = 0$ may mean that the individual is uninterested in participating and trading in a given economic sector, while $x_v = 1$ means the opposite; alternatively $x_v = 0$ could model that an individual does not have an active account for a social media service, with $x_v = 1$ the opposite. Each node will decide its state, 0 or 1, by comparing an internal field, which we label s_v, to an external (global) field p [27]:

$$x_v = \Theta(s_v - p). \tag{1}$$

For example, if we think of p as the external price of some good, then s_v represents the effective price at which buyer v is willing to buy: When $s_v > p$, $x_v = 1$ and the buyer is actively buying; when $s_v < p$, $x_v = 0$ and the buyer is not actively buying.

In order to fully specify the model we thus simply need to describe how to determine s_v. Our formulation of the binary decision model is to assume that

$$s_v \equiv P_v h(q_v), \tag{2}$$

where P_v is some internal variable, h is a monotonically increasing function, and

$$q_v = \mathbb{P}(x_u = 1|uv \in E) \tag{3}$$

simply represents the fraction of neighbors of v that are in state 1. Note that we will use $\mathbb{P}(\cdots)$ to denote probability throughout this paper. We are free to rescale P_v so that $h(0) = 1$ and we will assume so for the remainder of the paper. Note that we will always assume that h, P_v, and p are non-negative. Note that under the identifications

$$p = e^{p_+}, \tag{4a}$$

$$P_v = e^{P_{+v}}, \tag{4b}$$

$$h(q) = e^{h_+(q)}, \tag{4c}$$

we can formally write the binary decision model as

$$x_v = \Theta(h_+(q_v) + P_{+v} - p_+) \tag{5}$$

since the exponential is a monotonically increasing function.

To understand the choice above, it is helpful to recast the problem temporarily in the language of economics. Suppose that p represents some globally observed price (or, more abstractly, utility parameter); then P_v represents the price or utility that node v believes is the value of existing in state 1. The key point of this model is that the effective value of P_v is multiplied when neighbors of v are also in state 1. For example, a social media service (where P represents a utility, not a specific price) is far more valuable to its users when there are many other users. Similarly, stock traders may base much of their valuation of a stock based on what they believe the rest of the market is doing. Therefore, this binary decision model represents social scenarios where individuals are making a binary decision based on comparing the utility of two options, when the utility of one option is dependent on the behavior of their neighbors. Furthermore, the decision making is *cooperative* in the sense that if node v switches from $x_v = 0$ to 1, the probability that any neighbor of v is in state $x = 1$ is nondecreasing. This is in contrast to *antagonistic* decision making, where the probability that any neighbor of v is in state $x = 1$ is nonincreasing after v switches state. In this paper we consider only the cooperative case, where the individuals tend to make the same decisions as their neighbors.

The argument can and has been made [18] that an alternative choice

$$s_v = P_v + J \sum_{(uv)\in E} x_u, \tag{6}$$

which directly results in the famous random field Ising model, is also worth studying. In fact, note that this model is formally equivalent to the binary decision model on graphs where all nodes have the same number of edges for a special choice of $h_+(q) = kJq$. For the purposes of this paper, we will stick with Eq. (2), which leads to simpler calculations. We should note, however, that there is one major drawback to a choice such as this: As far as we can tell, there is in general no Hamiltonian function that will result in the binary decision model and this denies us the use of some of the tools of statistical mechanics.

III. MEAN FIELD THEORY

Let us now turn to mean field theory to solve the binary decision model and compare to simulations. Our first pass will demonstrate both how mean field theory can be quite accurate as well as reveal some of its major drawbacks.

A. Macroscopic solution

Let us define

$$q = \mathbb{P}(x_v = 1), \tag{7}$$

where the average over nodes is over the uniform distribution over V. Let us also assume that the random variables P_v are independent and identically distributed (i.i.d.), drawn from a probability distribution with cumulative distribution function (CDF) $1 - F(P)$ and probability density function $f(P)$. It is straightforward to see that

$$q = \mathbb{P}[h(q)P_v > p] = F\left(\frac{p}{h(q)}\right), \tag{8}$$

where the mean field approximation that $q_v = q$ has been applied. Equation (8) is the mean field equation describing the possible equilibria of the system, as described by the single parameter q. Note that this mean field equation appears to be independent of the degree distribution of the graph, if we choose to take this into account, since

$$q_k = \mathbb{P}(x_v = 1|k_v = k) = \mathbb{P}(x_v = 1) = q. \tag{9}$$

The middle equality above is a consequence of the fact that in mean field theory, every edge points to the same mean field node, i.e., every edge contributes the same factor of q. There are cases where this approximation will break down and we will return to this at the end of the paper.

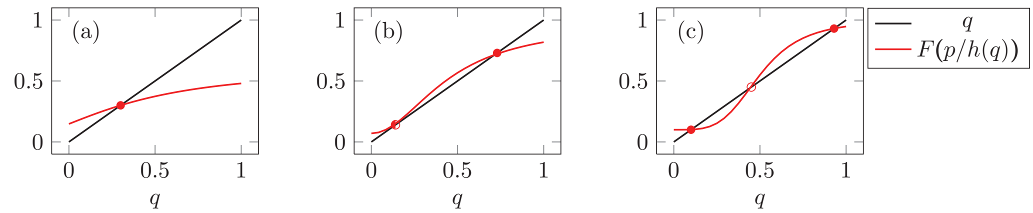

FIG. 1. (Color online) Graphical method of solving Eq. (8) in various cases: (a) a case with one solution, (b) a critical point, and (c) a bistable case with three solutions. The closed circles represent equilibria that are stable and the open circles represent equilibria that are unstable; the half-closed circle represents a marginal point.

We are most interested in the case where there are multiple equilibria, i.e., where Eq. (8) has multiple solutions. Graphically, it is clear how to approach this problem, as shown in Fig. 1. Inspired by this approach, we can also straightforwardly analyze it analytically. Since F is a CDF, we have $0 \leqslant F \leqslant 1$. Because we must have $0 \leqslant F(q=0)$ and $1 \geqslant F(q=1)$, there must be some q^* for which Eq. (8) is true. Furthermore, let us consider the behavior of F near this point. Suppose that $dF/dq > 1$, which means that $F < q$ for q just smaller than q^* and $F > q$ for q just larger than q^*. However, given the constraints at $q = 0$ and 1, we see that there must be at least two other points at which $q = F$, implying the existence of at least three solutions to Eq. (8). We conclude that if

$$\left.\frac{dF(p/h(q))}{dq}\right|_{q=q^*} = f\left(\frac{p}{h(q^*)}\right)\frac{ph'(q^*)}{h(q^*)^2} \equiv \alpha > 1, \quad (10)$$

then we have multiple equilibria. In a later section we will find a simple microscopic interpretation for α as the probability that a spontaneous flip in any node's state will cause one of its neighbors to also flip. The stability criterion $\alpha < 1$ follows from the condition that an avalanche have finite expected size.

Now let us consider what happens as we change p. In particular, suppose we increase p to some $p = p_c$ for which

$$f\left(\frac{p_c}{h(q^*)}\right)\frac{p_c h'(q^*)}{h(q^*)^2} = 1. \quad (11)$$

Then we conclude that passing through p_c in the direction that decreases the left-hand side of Eq. (11) corresponds to the disappearance of a pair of fixed points. In particular, if the solution at q^* disappears, it must discontinuously jump to another point. This is the hallmark of a discontinuous phase transition. It is not always the case that a distribution of $F(P)$ and an interaction $h(q)$ will allow such a discontinuous phase transition, but quite often discontinuous phase transitions do occur. Essentially, increasing the value of $h(q)$ makes the model more and more likely to admit a discontinuous phase transition and subsequently increase the size of the jump at the transition.

B. Exactly solvable case

Let us look at a sample of an exactly solvable version of this model to the level of approximation we have just studied. We will take

$$h(q) = 1 + Aq \quad (12)$$

and

$$F(P) = \begin{cases} 1, & P < P_0 \\ P_0 + 1 - P, & P_0 < P < P_0 + 1 \\ 0, & P > P_0 + 1. \end{cases} \quad (13)$$

It is easy to check when $q = 0$ is a solution: This occurs when $F(p/h(0)) = 0$ [or $F(p) = 0$] or $p \geqslant P_0 + 1$. Further, $q = 1$ is a solution if $F(p/(1+A)) = 1$ or $p < (1+A)P_0$. For $0 < q < 1$, solutions occur when

$$q = 1 + P_0 - \frac{p}{1 + Aq}, \quad (14)$$

which can easily be solved using the quadratic formula to give

$$q = \frac{1}{2}\left[P_0 + 1 - \frac{1}{A} \pm \sqrt{\left(P_0 + 1 + \frac{1}{A}\right)^2 - \frac{4p}{A}}\right]. \quad (15)$$

The important feature of Eq. (15) is the square root, which will become imaginary at price

$$p_c = \frac{A}{4}\left(P_0 + 1 + \frac{1}{A}\right)^2. \quad (16)$$

There are three distinct qualitative possibilities for the phase diagram. To understand these possibilities, it will suffice to consider q at $p = p_c$, which is

$$q_c = \frac{1}{2}\left(P_0 + 1 - \frac{1}{A}\right). \quad (17)$$

If $q_c > 1$, which occurs when $P_0 > 1 + A^{-1}$, then we know that only one branch of physical solutions exists (other than $q = 0$ or 1) and this branch has a positive slope in the pq plane, connecting $p = (1+A)P_0$ and $p = P_0 + 1$. If $q_c < 0$, then one branch of physical solutions exists, with negative slope, connecting the same two points. Otherwise, we see that q_c is a physical value and therefore is a valid point in parameter space; in particular, there are two branches of allowed solutions (other than $q = 0$ or 1). We show examples in Fig. 2.

Physically, we see that depending on the parameters, the binary decision model has a variety of interesting behaviors. When A is small enough, there are no discontinuous phase transitions, suggesting nice behavior. However, when A gets large enough, discontinuous transitions begin to occur and when A reaches a critical value, the only stable equilibria are with the entire graph in state 0 or 1, representing a hyperpolarized situation.

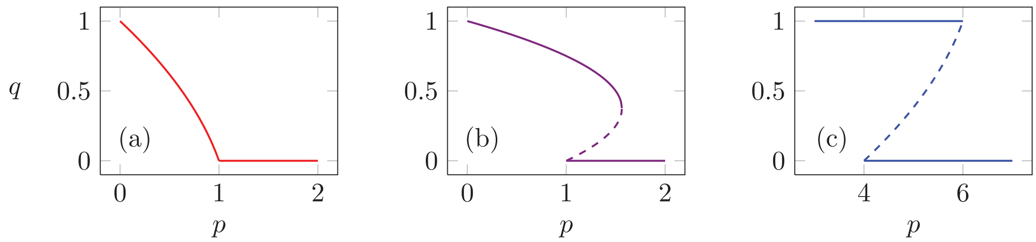

FIG. 2. (Color online) Three examples of exact mean field solutions for uniform distributions on P_v displaying qualitatively different behaviors: (a) $P_0 = 0$ and $A_0 = 0.5$, (b) $P_0 = 0$ and $A = 4$, and (c) $P_0 = 3$ and $A = 1$. The solid lines represent stable equilibria and the dashed lines unstable equilbria; we will understand stability from a microscopic standpoint later.

C. Numerical simulations

We now present numerical simulations of the binary decision model. In our numerical simulations, we always take $h(q)$ to have the linear form Eq. (12), although we allow the coefficient A to vary. We ran simulations over Erdős-Rényi random graphs, drawn from an ensemble where every possible edge is equally likely to be included in the graph. We then repeated simulations on asymptotically scale-free graphs, generated by a modification of the algorithm of Ref. [28] in which the added nodes have multiple edges. The algorithm is particularly easy to implement numerically, generates graphs that have degree distributions with robust scale-free ($\rho_k \sim k^{-\nu}$) tails, and can generate distributions for an arbitrary exponent $\nu > 2$. It does have the peculiarity that if each node that is added has m edges (in the large-N limit), $\langle k \rangle = 2m$, although this is just signifying that the small-k distribution is not scale free. Since we are interested in studying the model on scale-free graphs for any possible effects of nodes with very large k, this peculiarity is acceptable. We also note that many of our theoretical results are robust against the details of the graph ensemble and our simulations confirm this claim.

For our internal P distributions, we generated P_v from either uniform distributions over $[0,1]$ or Gaussian distributions with $\mu = 1, \sigma$ varying around 0.2, and scale-free distributions with varying exponents and minimum P_v of 1. We chose these distributions because they represent three different types of behavior, which could cause our mean field theory approximations to break down. The uniform distribution has substantial fractions of nodes with very low P_v; in contrast, the scale-free distributions have nodes with very large P_v. The Gaussian distribution tends to cluster nodes around $P_v \approx 1$. To observe bistability if it exists, we began our simulation by starting with $p = 0$, then increased p in uniform steps up until some value $p_{\max}$, and then decreased p in equal steps back to $p = 0$.

As a first simulation, we demonstrate that for a given internal disorder distribution $F(P)$ and a given $h(q)$, the mean field theory solution of Eq. (8) is typically a very good approximation. This is shown in Fig. 3. Note that, typically, the phase transition does not look as sharp as it occurred because we are averaging over runs and there are small fluctuations in the critical value of p due to internal disorder. We emphasize that the details of the graph ensemble, beyond $\langle k \rangle$, appear to have no effect on the curve $q(p)$. Figure 4 shows the emergence of a phase transition as the interaction strength A is increased past a critical value of 1 for the given uniform price distribution, as shown by both theory and numerics, as well as some sample distributions where P_v is a Gaussian or scale-free random variable. In each case, up to the deviations from mean field theory described above, we see excellent agreement with the theory, suggesting that for any P_v distribution, mean field theory will be a good approximation.

While the mean field theory approximation appears to typically hold, we should point out some *key* failures. In particular, there is a noticeable hysteresis effect associated with the upper branch of the mean field solution near the phase transition. We also notice that the phase transition appears to be slightly delayed for smaller values of $\langle k \rangle$ and on some graphs it appears as though there is no discontinuous phase transition at all. Furthermore, this effect cannot be removed by increasing N. Both of these phenomena are key signatures of the advertised multistability and we will return to them later.

D. Finite-size fluctuations from disorder

Let us briefly discuss the fluctuations about mean field theory due to finite-size effects (but *not* small $\langle k \rangle$ effects). These fluctuations are simply due to the fluctuations in the values of the internal fields P_v. To quantitatively estimate their

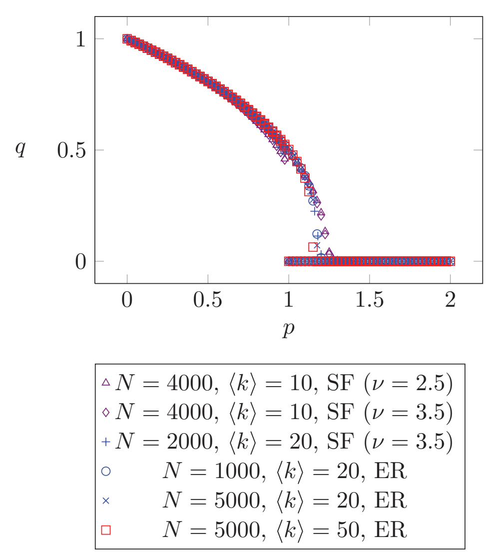

FIG. 3. (Color online) Average value of $q(p)$ for a uniform distribution of prices with $A = 2$. We averaged over at least 50 trials for each type of network. The smoother transitions correspond to graphs with smaller $\langle k \rangle$, a fact that we will explain later; note that there is only a single square ($\langle k \rangle = 50$) in the crossover regime, but many triangles ($\langle k \rangle = 10$). Here and in the following figures SF denotes the scale-free graph and ER denotes the Erdős-Rényi graph.

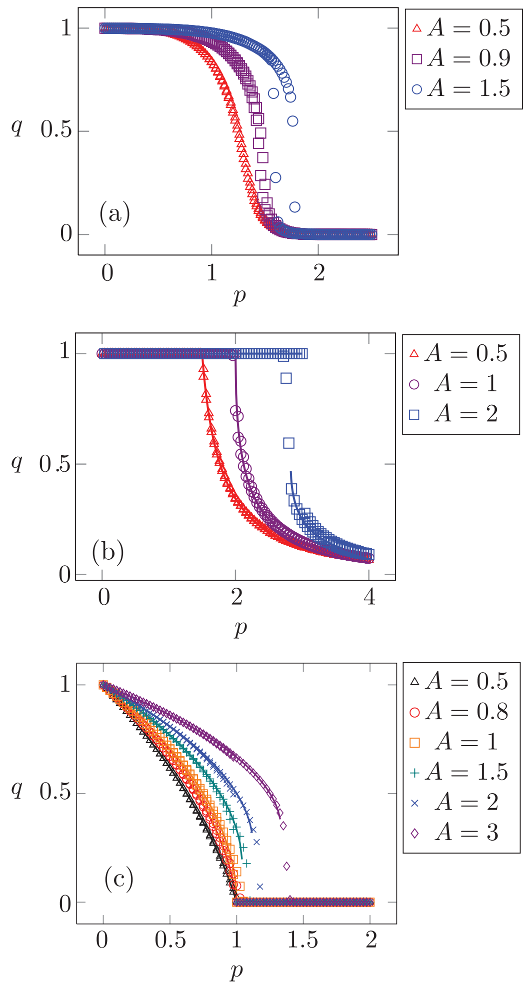

FIG. 4. (Color online) Here we show how increasing the interaction strength in the binary decision model causes the onset of a discontinuous phase transition in both theory and simulation. Our simulations for this graph used Erdős-Rényi graphs with 5000 nodes and $\langle k\rangle = 20$. The distribution of P_v is (a) Gaussian with mean 1 and standard deviation $\sigma = 0.3$, (b) scale free with exponent $\nu = 3$, and (c) uniform on [0,1]. Note that the discontinuous phase transition appears smeared out because of fluctuations in the realizations of disorder; data from individual runs clearly indicate the presence of discontinuous phase transitions for larger values of A.

size, let us define $q = q_0 + \delta$, with q_0 the mean field value predicted by Eq. (8) and δ a small fluctuation. Similarly, let us denote the CDF of the actual distribution of P, realized on the graph, by $F = F_0 + \Delta$, with F_0 the mean field value and Δ a small fluctuation. Then, expanding Eq. (8) to lowest order in the fluctuations we find

$$q_0 + \delta = F_0\left(\frac{p}{h(q_0+\delta)}\right) + \Delta\left(\frac{p}{h(q_0)}\right)$$
$$= q_0 + \alpha\delta + \Delta\left(\frac{p}{h(q_0)}\right), \quad (18)$$

which implies that

$$\delta = \frac{\Delta}{1-\alpha}. \quad (19)$$

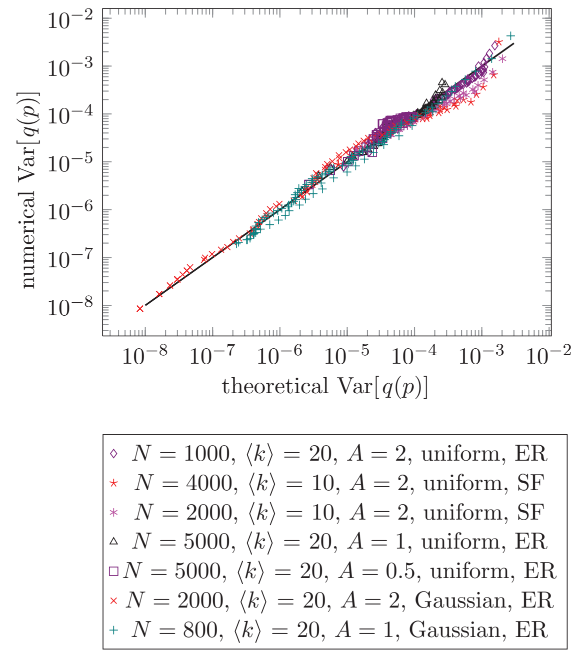

FIG. 5. (Color online) We compare the predictions of Eq. (21) to numerical simulations on a variety of graphs. The last two entries of the legend refer to the distribution of P_v and the graph type, respectively. For the data shown, Gaussian distributions have $\sigma = 0.3$ and scale-free graphs have $\nu = 2.5$. For clarity, only some of our data are shown, although we emphasize that the data not shown were just as good of a fit. Significant deviations come near the onset of a phase transition; we have removed some of these values near the critical point when our code is averaging over realizations in different phases.

The distribution of Δ is simple to find, although we will focus only on the variance of the fluctuations, as the higher-order fluctuations are rapidly suppressed for increasing N. Since the internal disorder consists of i.i.d. random variables and for the given value of p and the distribution of P_v the probability that a node is in state $x_v = 1$ is F_0, assuming the rest of the network to be in a mean field state, we conclude that

$$\mathrm{Var}(\Delta) = \mathrm{Var}\left(\frac{1}{N}\sum_v x_v\right) = \sum \frac{\mathrm{Var}(x_v)}{N^2}$$
$$= \frac{F_0(1-F_0)}{N} = \frac{q(1-q)}{N}. \quad (20)$$

This implies that

$$\mathrm{Var}(q) = \mathrm{Var}(\delta) = \frac{q(1-q)}{(1-\alpha)^2}\frac{1}{N}. \quad (21)$$

Since this number is typically quite small due to the factor of $1/N$ and our theory suggests that this variance is not dominated by large deviations, we postulate that higher-order fluctuations will not play an important role. We find that this relation is obeyed very well so long as we do not approach $\alpha = 1$, as shown in Fig. 5.

IV. AVALANCHES

We now turn to a slightly different question, which is a reinterpretation of our previous mean field analysis on the existence of a phase transition. Let us consider a state of the graph in equilibrium at $q^* = F(p_0/h(q^*))$. Suppose that the external field p_0 is increased slightly to p, causing only node v_0 to flip. How many other spins will also flip due to the change of state of v_0?

Let us define Q as the probability that one of the vertices v connected to v_0 will flip:

$$\begin{aligned} Q &= \mathbb{P}\left(\frac{p}{h(q_v - 1/k_v)} > P_v > \frac{p}{h(q_v)}\right) \\ &= F\left(\frac{p}{h(q)}\right) - \left\langle F\left(\frac{p}{h(q-1/k)}\right)\right\rangle_{\text{edges}} \\ &= F\left(\frac{p}{h(q)}\right) - F\left(\frac{p}{h(q)}\right) + \left\langle \frac{1}{k} f\left(\frac{p}{h(q)}\right)\frac{ph'(q)}{h(q)^2}\right\rangle_{\text{edges}} \\ &= \alpha\left\langle\frac{1}{k}\right\rangle_{\text{edges}}. \end{aligned} \quad (22)$$

Now we have to be a bit careful. As noted above, the averaging occurs over the distribution of nodes that an edge points to, which is different from the distribution of nodes itself because nodes with more edges are counted more often: In fact, each node v is counted k_v times in the average $\langle \cdots \rangle_{\text{edges}}$ over nodes pointed to by an edge. Therefore,

$$\left\langle\frac{1}{k}\right\rangle_{\text{edges}} = \sum \frac{k\rho_k}{\langle k\rangle}\frac{1}{k} = \frac{1}{\langle k\rangle}, \quad (23)$$

so we conclude that

$$Q = \frac{\alpha}{\langle k\rangle}. \quad (24)$$

Note that we can in fact make the stronger statement that the probability for one of v's neighbors to flip is independent of the degree of the neighbor.

Let us now denote by n the number of vertices v we expect to transition to the 0 state. For any given site, this occurs with probability Q, so

$$\langle n\rangle = \left\langle \sum_{j=1}^{k_{v_0}} Q\right\rangle = \sum \rho_k k Q = \langle k\rangle Q = \alpha. \quad (25)$$

Assuming that Q stays constant, if any new spin flips, then it too will have the possibility of flipping spins. If we approximate the new spin as the same as the old spin [29] then, since the avalanche approximately grows as a birth and death process, we conclude that the total size of the avalanche is

$$\langle n_{\text{total}}\rangle \approx \frac{\alpha}{1-\alpha}. \quad (26)$$

Of course, this formula holds only for $\alpha < 1$; for $\alpha > 1$ it is well known that the birth and death process has an expected infinite size and for $\alpha = 1$ it is almost surely finite but with infinite expected value.

This gives us much new insight into the physical processes at work behind mean field theory. Looking back at Eq. (11), we see that reaching a point where $\alpha = 1$ corresponds to a phase transition, where we expect a spin avalanche to have infinite size. Furthermore, we can now analyze the stability of the fixed points we found earlier: Only the fixed points where $\alpha < 1$ are stable. Understanding phase transitions by considering microscopic avalanche sizes is not new; see, e.g., Ref. [18] in the context of the random field Ising model or Refs. [20,21,30,31] in the context of the fiber bundle model. For our binary decision model, this calculation helps to give insight into why the graph structure is seemingly so unimportant in mean field theory. Even though more connected nodes are affected more often by spin flips, they are not affected as much and these effects cancel each other. This would *not* be the case in the random field Ising model on a heterogeneous network, for example, where α itself would obtain an intricate dependence on the degree distribution. In this sense, our binary decision model, where $h(q)$ is k independent, is a very convenient toy model to solve.

Let us also ask about the fluctuations in the size of avalanches, which we will explore by considering variations in the size of n, the number of neighbors that one site flipping will also flip. These fluctuations occur within the same realization of both graph structure and internal disorder. However, for simplicity we have included many realizations in our average. Defining $z_v = 1$ if v flips and 0 otherwise,

$$\begin{aligned} \text{Var}(n) &= \langle n^2\rangle - \langle n\rangle^2 \\ &= \sum \rho_k \left\langle\left(\sum_{j=1}^{k} z_j\right)^2\right\rangle - \alpha^2 \\ &= \sum \rho_k \left\langle \sum_{j=1}^{k} z_j^2 + \sum_{j\neq l} z_j z_l\right\rangle - \alpha^2 \\ &= \sum \rho_k [kQ + k(k-1)Q^2] - \alpha^2 \\ &= \alpha(1-\alpha) + \frac{\langle k^2 - k\rangle}{\langle k\rangle^2}\alpha^2. \end{aligned} \quad (27)$$

We see that the graph structure thus plays an important role in fluctuations of the size of avalanches. This theory is quite accurate as shown in Fig. 6, although it does seem to break down very close to a phase transition, as the diverging curves suggest.

Other work [30,31] considers in more detail the theory behind avalanche distributions in the fiber bundle model, which is quite similar to the binary decision model on fully connected networks. However, there is no general theory of the distribution of avalanches accounting for heterogeneous network structure and we have just seen how the variance of this distribution depends on details of the graph structure. Determining the full avalanche distribution, accounting for network structure, is likely intricate even in a mean field approximation.

V. MULTISTABILITY

In this section we discuss the binary decision model's most interesting feature: multistability. By the term multistability we mean that there is a continuous spectrum of equilibrium states, even in the $N \to \infty$ limit, so long as $\langle k\rangle$ is finite. Recall that an equilibrium for the binary decision model is a state where all nodes satisfy the constraint equation that $x_v = \Theta(P_v h(q_v) - p)$. Since q_v depends on the values of x_u for each neighbor u of node v, it is possible that there are multiple possible solutions to the constraint equation; thus

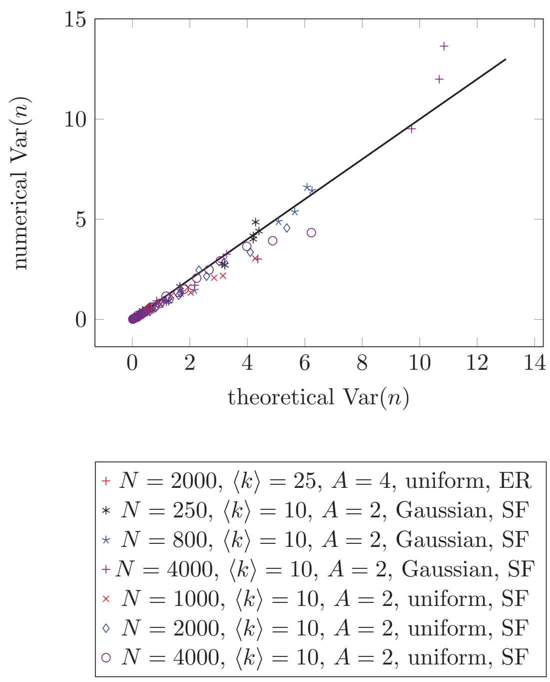

FIG. 6. (Color online) We compare theory vs numerics for the predicted variance in the number of nodes that change state based on the change of a single node. The deviations appear to become much more significant as we approach a phase transition, perhaps due to loopy effects in the graph (where our avalanche theory breaks down). In the data shown, scale-free graphs we used had $\nu = 2.5$ and Gaussian P_v distributions had $\sigma = 0.2$. We generated more data that appear very similar, but have not shown all of it for clarity.

multiple equilibria exist. We have trivially seen that this can be true by the existence of two phases, but here we are interested in the possibility that multiple equilibria may exist for any given phase.

Due to the internal disorder and the similarity to the random field Ising model, it is not surprising that glasslike behavior arises, although the lack of a formal Hamiltonian or free energy makes it challenging to classify the binary decision model as a glass. Instead, we will call the behavior multistability. The goal of this section is to propose multistability from a theoretical standpoint. We will then show it exists in our numerical simulations and comment on the implications of this phenomenon.

A. Spectrum of equilibria

Let us approximate the probability that a pair of spins will satisfy the x_v constraint if they are both in either state 1 or state 0. Note that it will never be the case that we would need to consider the possibility that a pair of nodes could flip between 10 and 01 because the constraint equation implies that decisions are always made cooperatively [32]. We will assume that the remainder of the graph is treated within the mean field approximation. The probability that one node flipping will flip another node has been calculated in the preceding section in Eq. (24). Thus we may calculate the probability that the pair can be in either state 1 or 0 by computing the probability node u flipping flips node v and vice versa:

$$\mathbb{P}(\text{pair is mutually 0 or 1}) = \left(\frac{\alpha}{\langle k\rangle}\right)^2. \tag{28}$$

Since there are $N\langle k\rangle/2$ edges in the graph, we thus approximate the number of equilibria as being

$$N_{\text{equilibria}} \sim 2^{N\alpha^2/2\langle k\rangle}. \tag{29}$$

However, we now need to take into account the expected value of the avalanche caused by the pair flip. This can be accounted for by roughly multiplying by a factor of $(1-\alpha)^{-1}$, corresponding to the expected size of the avalanche caused by one of the nodes. Thus we find that given a pair of nodes, we should expect that flipping them will cause a change of

$$\Delta q_{\text{one pair}} \sim \frac{1}{N}\frac{2\alpha^2}{(1-\alpha)\langle k\rangle^2}, \tag{30}$$

which leads to a width of the equilibrium spectrum that is finite even in the $N \to \infty$ limit, so long as $\langle k\rangle$ is finite:

$$\Delta q \sim \frac{\alpha^2}{(1-\alpha)\langle k\rangle}. \tag{31}$$

As discussed earlier, we ran simulations of the binary decision model by first increasing p from a state of all 1 and then decreasing p from states with many 0. This means that we can observe, quantitatively, the range of spectrum of equilibria by observing the difference between the value of q on the upward sweep versus the downward sweep. Figure 7 shows

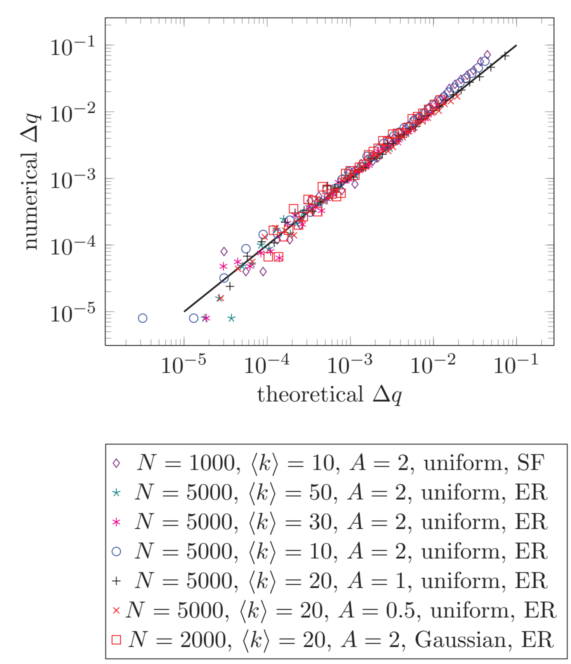

FIG. 7. (Color online) Comparison of the numerically determined width of the equilibria spectrum to Eq. (31). We see that the theory is very accurate even near a phase transition (these are the points with larger Δq). We have only shown a subset of the data for clarity. In the data shown, scale-free graphs have $\nu = 2.5$ and Gaussian P_v distributions have $\sigma = 0.2$.

that Eq. (31) is *quantitatively correct*, despite the incredibly simple theory we used.

Indeed, it should not be so surprising that we found characteristics of glassy behavior, considering the similarity of this model to the random field Ising model. Interestingly, it is unlikely this model is a true spin glass due to its similarity to the random field Ising model, where it was shown that the spin glass susceptibility is not divergent [33]. This is perhaps more of a mathematical technicality than a physically meaningful statement: In our model, so long as $\langle k\rangle$ is finite and $\alpha > 0$, the binary decision model is *always* a multistable glass.

B. Multistability by the cavity method

Since the multistability phenomenon is so fundamental to our model, let us predict multistability via a second approach: Thouless-Anderson-Palmer-like equations for x_v via the cavity method [34]. The cavity method is an extension of mean field theory that is used to understand disordered systems. We pick a special node in the graph, called the cavity node, and evaluate the probability that it is in each state, assuming that all of its neighbors feel the effect of the cavity's state exactly, with the remainder of the graph treated at a mean field level. Of course, this can be extended: An mth-order cavity method could take into account the reaction of the cavity on nodes up to m edges away. Note that the approach is also in the spirit of belief propagation on a tree [35].

To use the cavity method for our problem is straightforward: We need to make our mean field argument from earlier a bit more refined, so we will pick a special node v (the cavity) and explicitly write

$$x_v = \Theta\left(P_v - \frac{p}{h(q_v)}\right) \tag{32}$$

and explicitly determine q_v. To find q_v, we need to determine the expected value of a state x_u for each u with $uv \in E$. If we only use a first-order cavity approximation, then $k_u - 1$ of the edges of u point to nodes with $x = q$, so we find

$$x_u = \Theta\left(P_u - \frac{p}{h(q + (x_v - q)/k_u)}\right). \tag{33}$$

Averaging over the internal disorder *and* graph structure, it is straightforward to show that, assuming that k is large,

$$\begin{aligned}\mathbb{P}(x_u = 1) = q_v &= F\left(\frac{p}{h(q)}\right) + \frac{\alpha}{\langle k\rangle}(x_v - q)\\ &= q + \frac{\alpha}{\langle k\rangle}(x_v - q).\end{aligned} \tag{34}$$

Then we find the equation for the cavity node v, assuming again that k is large enough that we can Taylor expand $h(q)$:

$$x_v \approx \Theta\left(P_v - \frac{p}{h(q)} + \frac{ph(q)}{h'(q)^2}\frac{\alpha}{\langle k\rangle}(x_v - q)\right). \tag{35}$$

Now averaging over the internal disorder of P_v, we find that it may be possible that both $x_v = 0$ and 1 are solutions. The calculation of how likely this is very simple:

$$\begin{aligned}&\mathbb{P}\left(1 + \frac{h'(q)}{h(q)}\frac{\alpha q}{\langle k\rangle} > \frac{h(q)}{p}P_v > 1 + \frac{h'(q)}{h(q)}\frac{\alpha(q-1)}{\langle k\rangle}\right)\\ &\quad\approx f\left(\frac{p}{h(q)}\left(1 + \frac{h'(q)}{h(q)}\frac{\alpha q}{\langle k\rangle}\right)\right)\frac{\alpha}{\langle k\rangle}\frac{ph'(q)}{h(q)^2}\\ &\quad= \frac{\alpha^2}{\langle k\rangle} + O(\alpha^3),\end{aligned} \tag{36}$$

where the $O(\alpha^3)$ terms correspond to terms proportional to the derivative of f. This is precisely what we found earlier, neglecting the back reaction onto the remainder of the graph (via avalanches), up to the new terms that have arisen via the cavity method.

Let us now perform the cavity method to higher orders: In particular, let us assume that the graph is treelike (at least locally) and keep track of all the nodes up to a graph distance of n away from the cavity. The treelike approximation is very convenient as it allows us to assume that each node feels the effects of the cavity node v through exactly one neighbor. We start by considering the case of $n = 2$; the generalization to larger n will be very straightforward. The state of the cavity is given by

$$x_v = \Theta\left(P_v - \frac{p}{h\left(\frac{1}{k_v}\sum x_i\right)}\right) \tag{37}$$

and $x_1, \ldots, x_{k_v}$ denote the states of v's neighbors. The state of one of the neighbors of v, e.g., x_1, is given by

$$x_1 = \Theta\left(P_1 - \frac{p}{h\left(\frac{x_v}{k_1} + \left(1 - \frac{1}{k_1}\right)q_1\right)}\right), \tag{38}$$

where q_1 corresponds to the fraction of nodes (other than v) that are neighbors of node 1 and are in state 1. However, note that if x_1 is unknown, we already computed the formula for q_1 previously:

$$q_1 = q + \frac{\alpha}{\langle k\rangle}(x_1 - q), \tag{39}$$

where, as before, q is the fraction of nodes far from the graph. Now we write, assuming that the number of nodes is large, as usual,

$$\begin{aligned}x_1 = \Theta\Bigg(P_1 - \frac{p}{h(q)}\Bigg\{1 - \frac{h'(q)}{h(q)}\Bigg[\frac{\alpha}{\langle k\rangle}\left(1 - \frac{1}{k_1}\right)(x_1 - q)\\ + \frac{x_v}{k_1}\Bigg]\Bigg\}\Bigg)\end{aligned} \tag{40}$$

and using the same argument as before by finding the largest and smallest possible corrections to the effective value of p, we conclude that the probability that there is an equilibrium state with $x_1 = 1$, as well as one with $x_1 = 0$, given x_v, is given by

$$\mathbb{P}(x_1 = 0 \text{ or } 1|x_v) = \alpha\frac{\alpha}{\langle k\rangle}\left(1 - \frac{1}{k_1}\right) + O(\alpha^3). \tag{41}$$

Note that we have not yet allowed for fluctuations in x_0, which we have assumed is fixed. Also, given this formula, it should be fairly clear that we could have guessed this answer *a priori* from what we found before, under the assumption that one node x_v is fixed.

Finally, let us return to the question of interest: the probability that x_v can be in both states. We need to compute the largest and smallest possible values of q_v:

$$\mathbb{P}(x_1 = 1|x_v = 1) = q + \alpha\left[\frac{\alpha}{\langle k\rangle}\left(1 - \frac{1}{k_1}\right)(1-q) + \frac{1-q}{k_1}\right], \tag{42a}$$

$$\mathbb{P}(x_1 = 0|x_v = 0) = q + \alpha\left[\frac{\alpha}{\langle k\rangle}\left(1 - \frac{1}{k_1}\right)(-q) - \frac{q}{k_1}\right]. \tag{42b}$$

Note that the formula above neglects some of the smaller corrections due to derivatives in f, e.g., although these could easily be carried through if needed. We may finally average over the degree distribution and replace $1/k_1$ with $1/\langle k\rangle$. Using these equations to find upper and lower bounds on q_v, we finally can conclude, using the same logic as in our earlier computation, that

$$\mathbb{P}(x_v = 0 \text{ or } 1) = \alpha\left[\frac{\alpha}{\langle k\rangle} + \frac{\alpha^2}{\langle k\rangle}\left(1 - \frac{1}{\langle k\rangle}\right)\right] \approx \frac{\alpha^2(1+\alpha)}{\langle k\rangle}. \tag{43}$$

Note that this is also the width Δq of equilibria, up to second order in the cavity method.

Now let us extend this computation to higher orders. If we have found the width of the spectrum accounting for all nodes up to a distance n away, assuming the graph is a tree, we can easily determine the width of the spectrum accounting for all nodes a distance $n+1$ away from x_v by simply treating the neighbors of x_v as the cavities and using the result for n. An analogous computation to the one we did above shows that

$$\Delta q^{(n+1)} = \alpha\left[\frac{\alpha}{\langle k\rangle} + \Delta q^{(n)}\right]. \tag{44}$$

Using the results for $\Delta q^{(1)}$ and $\Delta q^{(2)}$ from before, it is clear that

$$\Delta q^{(n)} = \frac{1}{\langle k\rangle}\sum_{k=1}^{n}\alpha^{1+k}. \tag{45}$$

Summing this series to infinite order, we find that

$$\Delta q^{(\infty)} = \frac{\alpha^2}{(1-\alpha)\langle k\rangle}, \tag{46}$$

assuming that $\alpha < 1$ so that the series is summable. Thus we see that the cavity method recovers the result we found using simple logic earlier.

C. Robustness to fluctuations

An important consequence of the spectrum of equilibria is that the macroscopic model is *robust* against small perturbations in the external field p. To estimate the size ϵ of a perturbation in the external field p required to change the macroscopic state (the value of q) [36], we use the fact that

$$\frac{dq(p)}{dp} \sim \frac{1}{\epsilon}\frac{\alpha^2}{(1-\alpha)\langle k\rangle}. \tag{47}$$

Equation (47) follows directly from basic graphical considerations by considering Eq. (8) and using that if the slope of the mean field curve is known, then the fluctuations in both the q and p directions in $F(p/h(q))$ must be related. Taking an implicit derivative of Eq. (8), we find that

$$\frac{dq(p)}{dp} = -\frac{\alpha}{1-\alpha}\frac{h(q)}{ph'(q)}, \tag{48}$$

which implies that

$$\epsilon \sim \frac{\alpha}{\langle k\rangle}\frac{ph'(q)}{h(q)}. \tag{49}$$

For a generic model, the h-dependent factor in ϵ is likely $O(1)$ and so the dominant feature is the α and $\langle k\rangle$ dependence.

Equation (49) should be valid only away from a critical point since we made a linear approximation to relate its fluctuations in p and q. Near a critical point, nonlinear and nonanalytic effects in the relation between p and q become important. Indeed, Eq. (49) predicts that the binary decision model becomes more robust against small price fluctuations as we approach a critical point where $\alpha = 1$, a statement that is unlikely to be true.

D. Suppression of discontinuous phase transitions

What happens to the fluctuations of multistability as we approach the critical point? Here the simple linear response arguments we used above begin to fail. The $1/\langle k\rangle$ fluctuations can become so large on graphs with small values of $\langle k\rangle$ that they can in fact remove the discontinuity in the phase transition.

We can use a simple mean field theoretic argument to justify this. Let q_k denote the fraction of nodes with k edges in state 1 and r the probability that an edge points to a node in state 1:

$$r \equiv \frac{1}{\langle k\rangle}\sum_k k\rho_k q_k. \tag{50}$$

Accounting for possible fluctuations in the number of neighbors of each node that are in state 1, we can find a more precise expression for q_k than we gave in Eq. (9):

$$q_k = \sum_{m=0}^{k}\frac{k!}{m!(k-m)!}r^m(1-r)^m F\left(\frac{p}{h\left(\frac{m}{k}\right)}\right). \tag{51}$$

For simplicity, let us suppose that k is large enough that a binomial distribution is well approximated by a Gaussian distribution, so we have

$$q_k \approx \int_0^1 \frac{dq}{\sqrt{2\pi k^{-1}r(1-r)}}e^{-k(q-r)^2/2r(1-r)}F\left(\frac{p}{h(q)}\right). \tag{52}$$

Certainly when k gets very large, the Gaussian collapses to a δ function and we obtain $q_k = F(p/h(r))$, independently of k. However, suppose that the length scale q_F of the curvature (in the variable q) of $F(p/h(q))$ is comparable to $\sqrt{k^{-1}r(1-r)}$. Then we see above that fluctuations in the number of nodes that are actually in state 1 may average over a large enough region of q in $F(p/h(q))$ to remove the phase transition.

Since Eq. (51) is far too complicated to analyze in the regime of interest, where k becomes small, we resort to estimating numerically whether or not there will be a phase transition. We begin by determining self-consistent values of

TABLE I. The theoretical/numerical values for $\langle k\rangle$ at which we expect to see the onset of discontinuous phase transitions (into the $q = 0$ state) for the uniform distribution. Because these predictions require a precise understanding of multistability at the critical point, they should not be taken too seriously: We believe a standard deviation of ± 2 on each data point is not unreasonable (as we can only make scale-free graphs with even $\langle k\rangle$ with our algorithm).

Graph type	$A = 2$	$A = 3$	$A = 4$
scale-free, $\gamma = 2.5$	20/20	14/14	10/10
scale-free, $\gamma = 3$	18/20	12/14	10/10
scale-free, $\gamma = 4$	18/20	12/14	10/10
Erdős-Rényi	18/18	13/13	10/10

r, as calculated by using Eqs. (50) and (51). It is simplest to do this by plotting the function r and then the function determined by Eqs. (50) and (51) and looking for intersections. This graphical method also has the advantage that is easy to add in the effects of $1/\langle k\rangle$ fluctuations by simply adding $1/2\langle k\rangle$ to the right-hand side of Eq. (50). This follows from considering multistability as being caused by fluctuations of order $\alpha/2\langle k\rangle$: In fact, we can derive (31) in this way if we use the mean field expression for $F(p/h(q))$. The consequences of accounting for small $\langle k\rangle$ effects can become quite serious near a phase transition. Coupled with the effect of few edges smoothing out the weighting function $F(p/h(q))$ in the mean field equation, we find that theoretically we can see the disappearance of discontinuities in the phase transitions.

This gives us a crude way to estimate numerically whether a graph will admit a discontinuous or continuous phase transition. In Table I we perform the procedure above, assuming that $F(P)$ is the uniform distribution, and estimate the value of $\langle k\rangle$ for Erdős-Rényi or scale-free graphs for which we should see the onset of shocks and discontinuous transitions. A sample of what the adjusted mean field curve looks like, near the mean field critical point, with low $\langle k\rangle$ is also shown in Fig. 8. Numerical estimates suggest that this transition should be fairly sharp and we observed this numerically as well. Because we do not have a precise way of determining whether a phase transition has been discontinuous other than simply to observe the size of the change in q at each step forward in p, these results should be taken as at most semiquantitative. However, they do suggest that we have correctly identified the mechanism for the disappearance of discontinuous phase transitions.

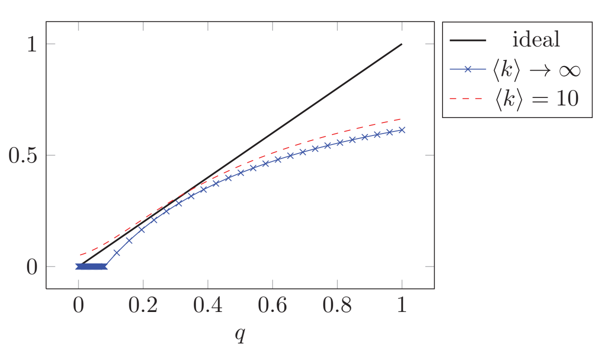

FIG. 8. (Color online) Sample of how small network effects alter the mean field equation, assuming that $F(P)$ corresponds to the uniform distribution, the network is Erdős-Rényi, and $h(q) = 1 + 2q$. The given value of $p = 1.16$ is quite close to the mean field critical point $p_c = 1.125$. We have added the $1/2\langle k\rangle$ correction to the finite $\langle k\rangle$ curve. Note that, while for the ideal case the only solution is $q = 0$, the small $\langle k\rangle$ curve still intersects the line at a positive value of q, i.e., the phase transition has not occurred yet.

Of course, there remains the question of whether or not these fluctuations actually merge the two phases together. This is a question beyond the scope of the simple arguments of this paper. We suggest that the answer may be that it depends on the network, the distribution of the disorder $F(P)$, and/or the interaction strength $h(q)$. For example, in the case with uniform $F(P)$, there is a very sharp discontinuous phase transition as the value of p is lowered from above to below the maximum allowed P_v. This transition is hard to remove because the coefficient of the $1/\langle k\rangle$ fluctuations is in fact 0 (as all nodes are in $x_v = 0$) and so network fluctuations will not affect this transition. Thus, at least for these systems, we conclude that there must be two distinct phases, although one of the transitions between the two phases may be continuous.

VI. CONCLUSION

We have presented the binary decision model and explored its major aspects: understanding from both a macroscopic and microscopic view the mean field solutions and fluctuations of basic quantities and then understanding the glassy multistability phenomena and the consequences of a spectrum of equilibria. The simple form of the decision making, as well as the interactions, allowed us to nearly exactly solve this model with very simple, physically motivated arguments. However, there are many questions about this model that are still open. First, it is unknown how robust the basic features of this model are to modifications: for example, heterogeneity among nodes in the interactions $h(q)$ or thermal random behavior among the nodes. Second, while it is unlikely that the multistability effect is a true spin glass effect, it is an interesting question whether or not this model's dynamics at finite temperature would exhibit aging, another characteristic feature of a glass. Third, an investigation of the model on nonrandom graphs such as hypercubic lattices or on graphs with many loops may help shed light on the nature of multistability, as we mentioned earlier.

It appears that on heterogeneous networks, the binary decision model has *fundamentally different* behavior from the random field Ising model despite the similarity in the motivation between the two models. One can intuitively see this as follows: For the random field Ising model with coupling J, using the additive formulation, the equation of state is

$$x_v = \Theta(P_v + Jk_v q_v - p). \quad (53)$$

This equation explicitly depends on k_v, the number of edges of v. Therefore, we expect the mean field equations to depend on graph structure, unlike in the binary decision model. Further study of qualitative differences in behavior between these two models, as well as modified fiber bundle models with the ability to regenerate nodes as p is decreased, is a worthwhile direction for further study.

Finally, it is important to understand the ways in which this model can be tested against empirical data. In general, this is quite challenging, but let us conclude by discussing some features of this model that may be observable. One of the most interesting features of our model is the remarkable robustness of our results against the degree distribution of the graph. For example, the probability that a node is involved in an avalanche is independent of its degree. Further, given $\langle k \rangle$, one could compare the width of the spectrum of equilibria to the typical size of avalanches and should find the same value of α from both measurements. There are two important drawbacks to this approach. First, trying to show the robustness of opinion dynamics against graph structure may require better knowledge of the social graph than can be obtained. Second, it is unclear how robust our formulas are to model modifications: For example, correlations between P_v and k_v may restore the degree distribution into many of our results.

The most important signature to look for in empirical data is multistability. This is unlikely to be a peculiarity of our model's precise formulation and discovery of such behavior would confirm the importance of emphasizing glassy features in opinion models. We leave further understanding of experimental signatures of this model, and related ones, to future work.

ACKNOWLEDGMENTS

We would like to thank Jean-Phillipe Bouchaud and Ariel Amir for helpful comments as well as the anonymous referees for suggestions. A.L. was supported by the Purcell Fellowship at Harvard. C.H.L. was supported by the Agency of Science, Technology and Research of Singapore.

[1] A. Barrat, M. Barthélemy, and A. Vespignani, *Dynamical Processes on Complex Networks* (Cambridge University Press, Cambridge, 2008).
[2] C. Castellano, S. Fortunato, and V. Loreto, Rev. Mod. Phys. **81**, 591 (2009).
[3] V. Sood and S. Redner, Phys. Rev. Lett. **94**, 178701 (2005).
[4] C. Castellano, M. A. Muñoz, and R. Pastor-Satorras, Phys. Rev. E **80**, 041129 (2009).
[5] D. Volovik and S. Redner, J. Stat. Mech. (2012) P04003.
[6] X. Castelló, A. Baronchelli, and V. Loreto, Eur. Phys. J. B **71**, 557 (2009).
[7] R. Axelrod, J. Conf. Res. **41**, 203 (1997).
[8] C. Castellano, M. Marsili, and A. Vespignani, Phys. Rev. Lett. **85**, 3536 (2000).
[9] F. Vazquez and S. Redner, Europhys. Lett. **78**, 18002 (2007).
[10] J. C. González-Avella, V. M. Eguíluz, M. G. Cosenza, K. Klemm, J. L. Herrera, and M. San Miguel, Phys. Rev. E **73**, 046119 (2006).
[11] S. Bornholdt, Int. J. Mod. Phys. C **12**, 667 (2001).
[12] E. Samanidou, E. Zschischang, D. Stauffer, and T. Lux, Rep. Prog. Phys. **70**, 409 (2007).
[13] D. Watts, Proc. Natl. Acad. Sci. USA **99**, 5766 (2002) .
[14] Y. Moreno, J. Gómen, and A. Pacheco, Europhys. Lett. **58**, 630 (2002).
[15] P. Crucitti, V. Latora, and M. Marchiori, Phys. Rev. E **69**, 045104(R) (2004).
[16] J. P. Sethna, K. Dahmen, S. Kartha, J. A. Krumhansl, B. W. Roberts, and J. D. Shore, Phys. Rev. Lett. **70**, 3347 (1993).
[17] S. Anderson, A. de Palma, and J. Thisse, *Discrete Choice Theory of Product Differentiation* (MIT Press, Cambridge, MA, 1992).
[18] J.-P. Bouchaud, J. Stat. Phys. (to be published).
[19] S. Galam and S. Moscovici, Eur. J. Soc. Psych. **21**, 49 (1991).
[20] S. Pradhan, A. Hansen, and B. Chakrabarti, Rev. Mod. Phys. **82**, 499 (2010).
[21] D.-H. Kim, B. J. Kim, and H. Jeong, Phys. Rev. Lett. **94**, 025501 (2005).
[22] J. Voit, *Statistical Mechanics of Financial Markets*, 3rd ed. (Springer, Berlin, 2002).
[23] S. Biswas, A. Chatterjee, and P. Sen, Physica A **391**, 3257 (2012).
[24] D.-H. Kim, G. J. Rodgers, B. Kahng, and D. Kim, Phys. Rev. E **71**, 056115 (2005).
[25] S. H. Lee, H. Jeong, and J. D. Noh, Phys. Rev. E **74**, 031118 (2006).
[26] H. G. Katzgraber, K. Janzen, and C. K. Thomas, Phys. Rev. E **86**, 031116 (2012).
[27] It is not important what happens when $s_v = p$ for almost all reasonable formulations of the graph G or the internal fields s_v.
[28] P. L. Krapivsky and S. Redner, Phys. Rev. E **63**, 066123 (2001).
[29] Arguably, the k above should be replaced with a $k - 1$, but for $k \gg 1$ this is not a major qualitative or quantitative change. We also require $k \gg 1$ for our Taylor approximation above to be accurate, so we will for simplicity neglect worrying about this for this paper.
[30] S. Zapperi, P. Ray, H. E. Stanley, and A. Vespignani, Phys. Rev. E **59**, 5049 (1999).
[31] M. Kloster, A. Hansen, and P. C. Hemmer, Phys. Rev. E **56**, 2615 (1997).
[32] With more work, this can be made rigorous.
[33] F. Krzakala, F. Ricci-Tersenghi, and L. Zdeborová, Phys. Rev. Lett. **104**, 207208 (2010).
[34] C. de Dominicis and I. Giardina, *Random Fields and Spin Glasses* (Cambridge University Press, Cambridge, 2006).
[35] M. Mézard and A. Montanari, *Information, Physics and Computation* (Oxford University Press, Oxford, 2009).
[36] It is possible that the equilibrium that we are at is at an outer edge of the spectrum, where this argument does not work, but for most equilibria (which are in the center of the spectrum) it works fine.