2 Modeling cognitions, networks, strategic games, and ecologies[1]

Matthew E. Brashears, Cornell University
David L. Sallach, Argonne National Laboratory

1. Introduction

Accurate knowledge of human behavior is necessary to deal with many strategic and tactical issues, but attaining such knowledge is difficult. Individuals are complex, self-reflective, and adaptive. Thus, a strategy that succeeds at one time may fail at another. Moreover, group behavior is not simply individual behavior scaled up; lessons learned at the individual level may not apply to a group and vice versa. In short, humans act and interact in complicated systems, and complex models are often necessary to describe them.

This chapter surveys some of the techniques that comprise the current state of the art in modeling human behavior at both the individual and group levels. The discussion is not exhaustive, as much work is being done in this area, but it provides a brief overview of some of the most promising methods currently in development. Specifically, we focus on various types of network-based models, cognitive models, and game-theoretical approaches. Network analysis constitutes an important and growing area of study and is vital in helping U.S. forces to confront decentralized opposition groups. Cognitive models have become essential due to the rise of self-starter terror cells, which plan and execute attacks with minimal or no coordination or resources from a central organization. Finally, game-theoretic approaches apply broadly to understanding equilibrium conditions across a wide range of levels of analysis, while promulgame theory concentrates on how game types interact to produce outcomes. To conserve space we have omitted other useful topics, such as agent-based simulation models, system dynamic models, and uncertainty analysis; these are covered elsewhere in this volume.

Computational models of human behavior differ in (a) the kinds of assumptions made about actors' knowledge; (b) the type of data required; (c) the representation of human activity; (d) the level of analysis; and (e) how adaptation is treated, if at all. We need a firm grasp of each of these factors to select an appropriate model and use it correctly.

Assumptions. All models impose some type of assumptions on the data, but those assumptions can differ markedly. For example, some models assume that individuals have perfect knowledge about

---

[1] Approved for Public Release; Distribution Unlimited. Case Number 14-2487
This work was supported by Office of Naval Research Contract N00014-10-C-0128

other group members, while other models assume limited knowledge. Empirical reality frequently violates some of the assumptions underlying a model, so users must assess whether the violations are serious enough to make the model unreliable.

Data Requirements. Models vary in the amount, and type, of data that they require as well as in their potential outputs. Some models require only very high-level demographic information, but produce high-level abstractions. Other models require detailed information on the characteristics and relationships of each modeled actor, but produce reliable estimates of individual attitudes and behaviors (e.g., technology adoption). Thus, both the outputs desired and the type and amount of data available drive model selection. However, there are no free lunches: the amount of data used to feed a model correlates at least roughly with the detail and accuracy of the model's output. This does not mean that accurate data always yields an accurate model, but only that it is difficult or impossible for a model to achieve a higher level of precision than is inherent in its inputs. Similarly, the data must accurately represent the process that analysts hope to explain. Biased or incomplete data will produce biased or incomplete results even when analyzed using the best available model.

Representation of Human Activity Change. Human activity has structure, but models differ in how they represent this structure. Some models assume that human relationships are relatively fixed, while others allow relationships to be born, mature, and finally die. The appropriate representation depends on the application; fixed relationships may be reasonable when modeling short spans of time, but not long spans of time.

Levels of Analysis. Models can be applied to levels of analysis that range from individuals, to small groups, to organizations, nations, and even alliances. Analysts should select models in part on the basis of the fit between the levels of analysis envisioned by the model and the particular entities studied. In many cases the same model may be appropriate for multiple levels of analysis, but the implications of model output can change in both obvious and subtle ways across these levels.

Adaptation. Humans are self-reflective entities, capable of understanding their surroundings and reacting to them. In other words, individuals or groups may recognize an intervention as an intervention and adjust their behavior or thinking in response. Computational models can represent these dynamics, but such modeling must be an ongoing process. Either a model of a social system should have reflexive capabilities, or analysts must update and revise it to take account of the changing empirical circumstances.

In the remainder of this chapter we explore several families of models useful for understanding social behavior, including Blau space ecological models, social network models, social influence models, cognitive models, game theory, and promulgame models. We begin with Blau space ecological models, which allow analysts to study network processes using demographic and membership data. These data are relatively easily obtained and thus Blau space approaches are useful for producing high-level abstractions. We then consider social network and social influence models, both of which allow for relatively precise estimation of tendencies toward particular relationship structures as well as behavioral and attitudinal predictions. However, both classes of models also require comparatively detailed data about individuals and their relationships and thus demand more data than Blau space models. Next, we address cognitive models and explore decision-making processes. These models focus on individual, rather than group, behavior and therefore can be relatively data intensive. Using these models with social network and influence models can yield valuable insights, but further increases the data required. Finally, we discuss game theory and promulgame models, which direct attention toward the individuals and groups evaluated in the earlier approaches, the features of the environment, and the interactions between them. Thus, the first main section of the chapter progresses from high-level abstractions through increasingly detailed and data-intensive models before concluding with approaches that deal with situational context.

We then describe the gaps and limitations affecting these models, thereby revealing some of the work that remains to be done. All of these models are subject to ongoing development and none has reached the level of turn-key solutions. Finally, we explore some of the steps that developers must take to transition these models into operational systems.

This chapter is intended to be useful to readers at a wide variety of levels of expertise, ranging from novice researchers to field practitioners. Needless to say, it is impossible to write a chapter that meets the information needs of all potential audiences, so portions of the content may be too technical for some readers and too simple for others. Thus, readers should take this chapter as an introduction that guides them through some of the most promising areas of research and gives a flavor of what these approaches can do. Sources listed in the bibliography provide more technical detail.

2. State of the Art

2.1. Blau Space Models

Blau space, named for Peter M. Blau, is a *k*-dimensional system in which each dimension is a sociodemographic variable and each point is represented by a set of *k* coordinates (McPherson, 1983; McPherson & Ranger-Moore, 1991). Each set of coordinates represents a unique combination of sociodemographic characteristics. Both individuals and organizations can be represented in Blau space at locations that correspond to these combinations. In other words, Blau space represents combinations of variables such as age, education level, and wealth as locations within a multidimensional space.

This multidimensional space conveys information because of homophily (Lazarsfeld & Merton, 1954), or the tendency for individuals to associate with those similar to themselves (e.g., “birds of a feather flock together”). Homophily is one of the most robust of all social science findings (e.g., McPherson, Smith-Lovin, & Cook, 2001) and plays a significant role in determining association in both the United States (Marsden, 1987, 1988) and elsewhere (Brashears, 2008). As a consequence of homophily, nearness in a Blau space model is positively related to the likelihood of association. Blau space therefore allows us to use demographic similarity to yield a probabilistic depiction of interpersonal ties, allowing analysts to produce an overall picture of population-level social networks without the need to collect expensive network data.

In addition to capturing individual networks, Blau space permits the modeling of ecological competition between organizations and cultural products (e.g., musical styles). All organizations require members and individuals are more likely to join an organization, or adopt a cultural product, if their associates have already done so. Even in the labor market network recruitment (e.g., referrals) accounts for approximately half of all hires (Bewley, 1999; Granovetter, 1995; Marsden & Gorman, 2001). Because individuals join groups to which their associates belong, and associates are typically homophilous (i.e., similar in multiple demographic attributes), organizations recruit from a limited and contiguous area of Blau space known as a "niche." Because both individual time and the number of persons in a given segment of a population (e.g., a niche), are limited, organizations whose niches overlap compete for the same finite pool of resources (i.e., member time). The larger the number of organizations that attempt to draw members from the same area in a Blau space model (i.e., the same population segment), the more competition these organizations experience.

The degree of competition in a region of Blau space has predictable effects on organizational memberships.[2] In areas of high competition, current members are more likely to leave the organization (McPherson & Rotolo, 1996; Popielarz & McPherson, 1995) and recruitment efforts will be less successful as rival organizations pursue the same individuals. In contrast, areas of relatively low competition produce more successful recruitment and less attrition. As a result, over time niches tend to move away from areas of high competition and towards areas of low competition.

Competing organizations often differ in their degree of specialization, which has important consequences for the competitive ecology. For example, a chess club may occupy only a small niche, while a church occupies a broad niche. To the extent that their niches overlap, these organizations compete for the time and attention of the same people. Competition is not symmetric, however, and the extent of the church's niche may mean that it only overlaps slightly with that of the chess club, whereas the church entirely encompasses the chess club's niche. Thus, the chess club competes with the church for all of its members while the church only competes with the chess club for a few of its members. A multi-organization system reaches equilibrium when the full carrying capacity (i.e., quantity of potential members) is allocated to the organizations in a manner determined by their degree of mutual competition.

While the organizational ecology can reach equilibrium, periodic changes in the carrying capacity generate shifts in the positions of organization niches and prevent a static equilibrium (McPherson, 2004). However, the direction of niche movement can be predicted, allowing analysts to assess the future composition of a group (Bonikowski, 2010; McPherson & Ranger-Moore, 1991). Thus, Blau space permits us to construct an ecological model of organizational membership and to predict the changing demographics of organizations—at least in part—through a competitive process of recruitment and retention.

---

[2] Organizations need not realize that they are competing. As long as two organizations exploit the same finite pool of resources, they are in fact competing.

The affiliational ecology model was originally constructed to explain the membership characteristics of organizations, but has also been applied to cultural products (Mark, 2003) such as musical styles (Mark, 1998), including change over time (Bonikowski, 2010). In other words, cultural products, such as musical styles, can be modeled as competing for supporters in much the same way that organizations compete for members. In theory, a wide variety of phenomena can be modeled in a similar fashion (McPherson, 2004).

McPherson (1983) first proposed Blau space models, and this approach continues to develop. Recent research, funded by the Defense Threat Reduction Agency, has continued to elaborate and refine the basic model. This stream of research has shown that models using Blau space methods can predict both deviant and conventional behaviors at the individual level (Brashears, Genkin, & Suh, 2012; Suh, Brashears, & Genkin, 2012). Importantly, organizational competition also influences the behavior of organization non-members (Brashears, Genkin, & Suh 2012; Suh, Brashears, & Genkin, 2012). In combination with earlier work, this suggests the possibility of predicting individuals' tendencies to engage in various types of behaviors (e.g., suicide bombing) using an inexpensive Blau space approach.

Additional research seeks to generalize the concept of the niche to individuals. In this model, an individual's Blau space neighbors represent a set of others who have regular contact with that person and thus exert consistent influence, regardless of whether they are directly connected to the individual. Such others often share a physical or social context with the individual, or are connected via third parties, and thus serve as both a reference group and a measure of the local informational environment. This method has proven effective in predicting health perceptions and behaviors in a U.S. sample, even when controlling for a wide variety of covariates and direct network effects (Behler, Suh, Brashears, & Shi, 2013). Thus, Blau space models offer a powerful and flexible way to model human behavior using minimal data, but generally produce a high-level abstraction of the modeled social system.

2.2. Models for Social Networks

Human social networks are of increasing interest to scholars and analysts, but drawing valid statistical inferences about these networks poses challenges. Most statistical techniques assume that observations are independent (e.g., the opinion of one respondent to a survey does not determine the opinion of another respondent), but this assumption cannot apply to social networks. Individuals who share a relationship (i.e., are tied) influence one another, and thus standard statistical methods cannot be employed (e.g., Marsden & Friedkin, 1993). Additionally, network data are expensive to collect and few analysts have access to a "sample" of networks. As a result, analysts cannot use standard techniques to compute statistically valid estimates of network-level processes. Both of these difficulties are exacerbated when dealing with data collected over time, which allows more certain identification of causal relationships, but also introduces statistical dependence on a longitudinal as well as a cross-sectional basis. Finally, analysts often wish to explore how network structures (e.g., who is friends with whom) and actor attributes (e.g., characteristics and behaviors) influence each other, thereby creating even more complex dependency structures. For example, are individuals influenced to become suicide bombers by their associates, or do individuals who want to become suicide bombers tend to become associates? Fitting models to network data calls for specialized techniques, and two types of models have gained considerable traction: Exponential Random Graph Models (ERGMs), which parameterize graph features, and Siena models, which parameterize the tendencies of simulated actors.

ERGMs begin with the observation that any realized network is a specific instance drawn from a population of potential networks. For example, a group of 15 individuals can have a total of $n(n–1) = 15(15–1) = 210$ relationships. If these relationships are binary (i.e., a relationship either exists or does not exist), then the sample space from which the observed network is drawn consists of 2^{210} or 1.64×10^{63} possible networks. The probability of observing any specific network in this space, assuming all are equally likely, is $1/2^{210}$ (Wasserman & Faust, 1997). However, while all networks might be equally likely to occur, certain network characteristics are more likely to occur than others. For example, only one network contains all 210 possible ties and only one network contains 0 ties, but many networks contain 105 ties.[3] Thus, models can generate expectations for how many ties a network drawn at random from this sample space should contain. If the observed network contains many more or many fewer ties than this expected value, we can infer that some mechanism is encouraging, or discouraging, the formation of social ties. The same logic extends to other network processes, including reciprocity (i.e., the tendency for individuals to select each other as associates) and transitivity (i.e., the friend of my friend is my friend), among others. ERGMs estimate parameters describing how much more (or less) often a particular feature is observed than we would expect by chance.

The statements above require two elaborations. First, because many characteristics of a network become more probable when many ties exist (e.g., reciprocity is higher when there are more ties overall), parameter estimation must control for the number of ties in the network. ERGMs impose this control by limiting the sample space to those graphs with the same number of nodes (i.e., vertices) and ties (i.e., edges), as the observed graph. Second, while the ties are binary, the probability that a tie exists does not have to be 0.50. Applying a Bernoulli distribution can model the probability that a tie exists across the potential range, and these probabilities can be conditioned on other factors, including the size and density of the network as well as the attributes of actors and actor pairs, yielding a set of conditional uniform distributions.

The general ERG model (Lusher, Koskinen, & Robins, 2013; Robins, Pattison, Kalish, & Lusher, 2007) parameterizes particular network features. One feature might be a closed triad (i.e., a group of three nodes that all have connections to each other) while another might be an open triad (i.e., a group of three nodes that contains only two ties). The model calculates statistical coefficients representing the tendency for a given graph to contain more, or fewer, of these features than would be expected by chance. It determines parameters using Monte Carlo Markov Chain Maximum Likelihood Estimation techniques, which permit inclusion of a very large variety of

[3] To use a more familiar example, when rolling two (fair) six-sided dice, only one combination sums to 2 (1+1) and one combination sums to 12 (6+6), but six combinations sum to 7 (1+6, 6+1, 2+5, 5+2, 3+4, 4+3). Thus, even though all possible combinations are equally likely, combinations that sum to 7 are much more likely to occur than combinations summing to 2 or 12.

configurations (Snijders, Pattison, Robins, & Handcock, 2006; Wasserman & Robins, 2005), although estimating parameters with larger numbers of nodes and ties can consume considerable time. These techniques generate a set of graphs using randomized starting values for the model parameters and compare them to the observed graph; if the observed graph has a very low likelihood of belonging to this set, the model selects new parameter values and repeats the process. Over many replications, the process homes in on the parameter values most likely to have produced the observed graph. Modelers must exercise care, however, to ensure that parameter estimates remain stable against local maxima (i.e., models should be estimated several times with different initial seeds). They may also have to include some configurations in the model regardless of theory in order to avoid near degeneracy, or a situation in which only a small subset of all possible graphs has other than a very low probability (Handcock, 2003). A good fitting model can be defined as one that yields a very high likelihood of obtaining the observed graph relative to the alternatives.

Scholars continue to develop the ERGM framework. Wang, Robins, Pattison, and Lazega (2013) have generalized ERGMs to accommodate two-level scenarios (e.g., individuals nested within groups). This approach resembles hierarchical linear modeling and allows analysts to examine interactions between individual relations and higher level structures (e.g., organizations). However, the software implementation of this model is currently somewhat unreliable. Other work (Caimo & Friel, 2012) has wedded ERGMs to a Bayesian framework to assess the quality of model fit. This helps to avoid reliance on frequentist model fit measures, which depend upon as yet poorly understood comparison distributions. Thus, in addition to considerable empirical work using ERGMs, the method continues to acquire methodological sophistication.

ERGMs capture social processes by parameterizing features of the network, but impose no assumptions about the mechanisms that generate these features. This adds to the model's flexibility, but can leave the analyst in the position of trying to justify configurations that lack a clear theoretical meaning. Stochastic actor-based models for network dynamics (i.e., Siena models) avoid this problem by building stronger assumptions into the model. Additionally, Siena models are ideal for longitudinal data.

A Siena model represents interpersonal ties as states rather than events, meaning that it shows any particular communication event (e.g., an email exchange) simply as an expression of an underlying, ongoing relationship. States can change (i.e., ties form and are dissolved) but in general existing ties tend to continue existing and nonexistent ties tend to continue not existing. Time varies continuously in arbitrarily defined segments and actors control their outgoing ties (i.e., individuals choose whom they prefer) but not their incoming ties (i.e., individuals cannot choose who prefers them). Finally, Siena models employ simulated actors, or simplified mathematical representations of the entities (e.g., individuals) that make up the nodes of the network.

A Siena model proceeds via two mechanisms: the change opportunity process and the change determination process. The change opportunity process determines the rate at which actors in the model have the opportunity to alter their network ties or behaviors. In its simplest form, this rate is uniform across all actors, meaning that at any given time all actors are equally likely to make a discrete change (i.e., all individuals change at the same rate). Alternatively, the rate can be conditional, meaning that some actors (e.g., more popular individuals) are more likely to make discrete changes than others (i.e., exhibit a higher rate of change). The rate of change can also vary over time (e.g., change occurs at an increasing rate). The model determines stochastically which actor will make a change, with the rate controlling the probability that any given actor will be chosen.

The change determination process models the changes that the actors make when given the chance. Actors may make a single change per opportunity: add a single tie, eliminate a single tie, change one attribute, or do nothing. This one-change rule defines a finite set of alternative networks that actors can reach at the present time. The actor's *objective function*, which defines how rewarding particular network structures are for the individual actor, determines which change is made (i.e., which alternative network is selected). Put simply, the actor will tend to select the alternative that produces the most positive objective function. However, because the model is stochastic, actors do not always select the change that yields the greatest improvement in the objective function. Thus, while the objective function can be thought of in terms of "preferences," it is more accurate to view it as capturing tendencies stemming from preferences as well as circumstances beyond individual control (e.g., the availability of certain types of others in the population).

The effects in a Siena model refer to the tendencies of simulated actors to make specific types of changes (Snijders, Van de Bunt, & Steglich, 2010). The effects can be defined by structure (e.g., reciprocity, transitivity), by attributes (e.g., level of behavior adoption), as well as interactions between the two (e.g., homophily).

The Siena approach uses a method-of-moments procedure to generate parameter estimates and therefore models the process of change between time points rather than tendencies at any one time. The first time point is not modeled, serving instead as a start for the estimation process. As a result, a Siena model cannot be estimated without data from at least two points in time.

In outline, the Siena model proceeds by taking the first wave of data (t) as the starting point for the simulation model. An actor is chosen at random, with the probability determined by the parameters of the change opportunity rate (uniform or conditional). The model computes the objective function for both their current network structure as well as all network structures reachable via only a single change, with the probability of a given change determined by the size of each objective function, including the function for the existing network structure. Once the model has selected a given network using this procedure, the change (if any) is made and a new actor is chosen via the change opportunity process. Change in the network is modeled as a Markov process: the state of the network at the current time (t) probabilistically determines the future (t+1) state of the network, with no additional effects felt from prior (t–k) states. This process repeats many times, in theory yielding outcome networks comparable to the observed networks at later time points (i.e., t+1, t+2, etc.).

Initially the analyst sets the parameters in the objective function randomly and the software compares the performance of the simulation to the observed data. If the simulation fares poorly the values are adjusted and the simulation process is repeated. Eventually, the model homes in on the best set of parameter values (including those for the rate of change parameters) for reproducing the observed networks. As with ERGMs, however, modelers must take care to ensure that parameters have not been trapped in local maxima; re-estimating with new seed values is advised.

Recent elaborations of the Siena model have explored its applicability to two-mode data (Snijders et al., 2012), allowing researchers to connect individual and organizational processes in a single longitudinal model. Brandes, Indlekofer, and Mader (2012) have also developed diagnostic visualization tools that assist researchers in determining the quality of model fit. Thus, the Siena approach is still undergoing considerable development.

Siena models and ERGMs differ in their overall perspective on social networks. Siena models view the actors as making changes to their networks in pursuit of particular structures. This has the advantage of producing parameters that allow many small changes to a network to accumulate into overall network configurations, but runs the risk of treating actors as having more control over the social system than they really do. ERGMs, in contrast, focus on tendencies for particular features to appear more often than expected by chance, but make no effort to model the processes that give rise to these features. Therefore, analysts can identify features that are preferred more often than others, but that cannot be realized through any straightforward process. The model preferred for a given application should depend on data availability, on whether the data are longitudinal or cross-sectional, and on whether the nodes in the network can reasonably be thought of as “agents” (e.g., transformers in a power grid are not “agents” in the same sense as human individuals in a social network).

2.3. Social Influence Network Theory

To understand the hearts and minds of a population we must study attitudes. Social Influence Network Theory (SINT) fills this gap by modeling valued attitudes, meaning affective (i.e., emotional) views of some issue that can range from very positive to very negative. The theory assumes that individuals enter an interaction with some initial attitude and are susceptible to influence from others. The degree of influence is determined by three characteristics: the network of relations among individuals, the weights placed on these relations, and the attitude of each relationship partner.

First, individuals are tied to others in a social network. We exchange information, views, and camaraderie with our associates, and their attitudes can influence our own views. Additionally, however, our associates have relationships to others to whom we are not connected (i.e., our second-order network). As those others influence the attitudes of our associates they will, indirectly, influence our own attitudes. To understand the process of attitude change and consensus, we must model the entire network.

Second, we do not necessarily listen to each of our associates equally carefully. Some of them have a strong influence on our own views, while we may tend to ignore others. As a result, the model must take account of the unequal influence of various associates, and must factor in our relative level of self-confidence.

Finally, the extent to which initial attitudes must change in order to reach consensus is at least somewhat a function of starting positions. Individuals who begin relatively further apart on an issue must change more in order to reach consensus than those who begin near each other.

The combination of these straightforward elements yields complex and interesting behavior, as each can moderate or enhance the impact of the others. For example, very different initial positions may result in relatively little attitude change if few direct network connections exist between opposite extremes and if the self-weights are very high. Alternatively, when others are very influential, attitudes may change so rapidly that the system will repeatedly overshoot consensus, ultimately requiring longer to reach equilibrium. Therefore, while the model is conceptually simple, it has rich implications.

The SINT model (Friedkin & Johnsen, 2011) includes several important elements. First, an actor's self-confidence is locked in a zero-sum game with the influence of others; actors less certain of their own views are more susceptible to the attitudes of others. Second, all persons to whom an actor is tied have an influence on the actor's future views, with the degree of influence determined by the product of their level of influence and their current attitude. Third, the views of the actor at the first time point (t=1) exert a continuing influence in the model at all points in time. This last element is extremely useful in preventing the model from predicting an unrealistic degree of consensus, but also treats the first time point as special relative to the remaining points. In fact, because the analyst chooses the first time point based on convenience or data availability, its theoretical value is dubious. Finally, the model is computed iteratively; each computed set of attitudes is reentered to predict the attitudes at the next time step until the system reaches equilibrium.

SINT allows an analyst to predict the equilibrium states of members, but two or more consensus positions are possible at equilibrium, implying some ongoing disagreement. In all but a very few boundary cases the equilibrium consensus positions lie on the convex hull of the initial positions.[4] Put differently, the equilibrium consensus will not be more extreme along any attitude dimension than the most extreme initial positions in those dimensions. The model has an additional advantage in that it can derive many of the quantities needed from partial information. For example, if data on initial and ending attitudes are available, as well as network data on the actors,

[4] In dyads where both partners heavily influence the other, consensus positions sometimes, though rarely, appear outside the convex hull. This appears to result from an attempt to escape from a coordination problem. In these systems interaction partners repeatedly overshoot each other in an attempt to reach agreement, and so one partner adopts a more extreme position in order to provide a conspicuous target for the other. For example, a married couple might try to choose a restaurant for dinner. Each suggests a restaurant and each attempts to defer to the other's preference, thus preventing consensus and delaying the meal. One partner might attempt to break the cycle by suggesting a third option that usually neither prefers. While this example is only loosely accurate (i.e., the example deals with a decision rather than an attitude), it does help illustrate the general process.

the model can derive the weights on those relationships. Thus SINT provides a flexible and effective way to model attitude change and stability in groups.

2.4. Cognitive Models

Researchers have used a variety of mathematical and computational models to emulate the cognitive process. The examples summarized below highlight some of the most influential.

An early example focused on Belief, Desire, and Intention (BDI). Emerging from artificial intelligence, the BDI model treats the three constitutive concepts as key pragmatic dimensions of cognition (Bratman, 1987). Belief encompasses those assertions that the actor considers true and that, in some cases, act as a placeholder for the actor ontology (although the actor may not be consciously aware of them). Desire conveys focus and is sometimes used interchangeably with goals. Intention describes the phase in which belief and desire have congealed into a settled plan, or course of action.

BDI models haves been criticized for doing both too much and too little. For example, some critics say that having incorporating only three steps makes the mechanism underlying BDI models too simple, while others view three interacting dimensions as overkill (Rao & Georgeff, 1995). However, BDI models suffer from their ultimate limitation in that, while seemingly intuitive, they have no obvious foundation in either decision theory or substantive social action that would shape and drive their computational implementation and/or the resulting dynamics. Such criticism suggests that BDI is constructed on the basis of folk psychology and, thus, cannot provide models with more than a pseudo-cognitive placeholder.

The rational choice model of cognition, as developed by Becker (1976), Coleman (1990), and numerous others (cf., Elster 1986), has played an influential role in the social sciences, albeit arousing controversies along the way (Coleman & Fararo, 1992; Green & Shapiro, 1994). The formulation of individual models depends on the domain and the research strategy, but all rational choice models provide a mechanism for assessing the costs, benefits, and tradeoffs of different courses of action, and the responses that result from such assessments.

One strength of rational choice lies in its high level of generalizability. Researchers have used the paradigm in a wide range of social domains, from microeconomics (Field, 1984) to international relations (Levy, 1997). Often researchers have been able to identify rationalities underlying strategies that, on the surface, appeared to be irrational or self-defeating (Pape, 2005).

However, a primary weakness of rational choice is that it typically involves closed concepts of gain and loss. Modelers often act as if these concepts are self-evident and, as a result, many models weakly represent the dynamic spectrum of priorities that social actors embrace, including their evolution over time. Another problematic characteristic concerns how the models represent emotion as a factor in decision processes. Often they simply ignore affect and, when they do address emotions, frequently view them primarily as a source of irrationality (Doran, 2000). This undermines possible analysis of strongly committed action and/or the complementary ways in which emotion and cognition work to achieve insights and effective action.

Randall Collins has been a major contributor to cognitive models in the social sciences that integrate both emotion and ritual, creating a more extended concept of cognition. He developed his theory of emotion in social life over a period of time (Collins, 1981a, 1981b, 1990; Kemper & Collins, 1990), ultimately proposing a mechanism designed to subsume rational choice models (1993).

Collins' incorporation of emotion into sociological theory is based upon Durkheim's discussion of religious ritual (1995) and Goffman's subsequent analysis of interaction rituals (1967), as unified in his own theory of interaction ritual chains (IRCs) (Collins, 1981a, 1981b, 1990, 2005). In brief, Collins argues that group solidarity at all levels is grounded in emotion, that emotional bonds precede and make possible coercive and contractual relations, and that IRCs of widely varying types generate group solidarity as well as emotional energy (EE) among the participants. The efficacy of specific IRCs varies depending on the attributes of group interaction: physical assembly, boundedness, focus of attention and awareness, commonality of emotional mood, and symbols that represent group membership (Collins, 1988; Collins, 1993).

Collins (1993) presented IRC/EE as a common metric and motivation underlying all social exchanges. His model recognizes “emotional solidarity with a group as the primary good in social interaction” (1993, p. 205). He notes that individuals “are motivated to maximize the amount of solidarity they can receive, relative to the costs of producing it” (1993, p. 209).

Nonetheless, Collins's IRC/EE model has weaknesses that prevent realization of its putative advantages. First, contrary to Collins's assertions, there is little reason to believe that the interchange of interaction rituals fully determines economic markets. Second, serious questions arise as to whether IRC markets operate in the narrow way that the model suggests. However, Collins's model has served as a springboard to a more general affect-based field theory (Sallach, 2008) that avoids the limitations of the IRC/EE model in its speculative 1993 formulation while, at the same time, integrating emotional and cognitive processes for social actors at multiple scales. This integration allows behavior previously regarded as generically “irrational” to more specifically recognize the underlying emotional commitments that override utilitarian considerations, sometimes in a strongly persistent manner (cf., Benmelech & Berrebi, 2007).

2.5. Game Theoretic Models

The formalisms most frequently associated with adversarial interaction derive from game theory. They have been applied widely, and are one of the rare forms of social analysis that can explore decision processes and outcomes using deductive analysis, while increasingly supporting advanced simulation environments as well (cf., Parsons, Gmytrasiewicz, & Wooldridge, 2002). Fudenberg and Tirole (1987) summarize the several sources of game-theoretical strength:

> [Game theory] imposes some discipline on theoretical thinking. It forces economists to clearly specify the strategic variables, their timing and the information structure faced by firms . . . [T]he researcher learns as much from constructing the model (the "extensive form") as from solving it because in constructing the model one is lead to examine its realism. (p. 176)

However, it also appears that the more complex and dynamic the social phenomena under study, the more difficult it is to achieve effective game-theoretic representation. Peltzman (1991), for example, bemoans the seeming inability of theory, especially game theory, to lead to powerful generalizations. He finds, rather, " . . . an almost interminable series of special cases. The conclusions drawn from these cases tend to be very sensitive to the way problems are defined and to the assumptions that follow" (p. 206). Peltzman continues: "By suitably permuting and combining the problems and assumptions, new models can be produced almost *ad libitum.* Indeed, the production of new models and the tidying up of old ones seem to be major goals of this research enterprise" (p. 207). Similarly, Fisher (1989) writes: "There is a strong tendency for even the best practitioners to concentrate on analytically interesting questions rather than on ones that really matter for the study of real-life industries" (p. 123).

Game theory has provided a mathematical way of exploring interactive social outcomes under varying assumptions. However, its very formalism has led to concerns about unrealistic assumptions (e.g., perfect foresight, stylized outcomes) and a failure to incorporate the situated contexts that inform strategic decisions. To maximize its theoretical and practical contributions to social modeling, game-theoretic models need to become more dynamic. The resulting models should be elegant, and guided by social theory and, as a result, they will also be considerably more expressive.

Human actors play multiple games, select among available games, shift from one to another, sometimes mistake what the game their counterpart is playing, act in ways that are (more or less) effective within multiple games simultaneously, etc. Luttwak (1987) points out that strategic choices are inherently paradoxical: because the strongest, most effective options are also the most obvious and, thus, are likely to be the best defended. Options that are suboptimal, and thus unexpected, are sometimes the most effective. This emphasis upon surprise is not just a product of strategic planning. As Miller (1997) has documented, unpredictability has been essential to species survival and adaptive dominance. Accordingly, interpretive dynamics that produce unexpected results can be seen to be essential to the accomplishment of participatory and analytical understandings of strategies and courses of action.

2.6. Promulgames

Promulgame theory represents a cluster of innovations within game theory designed to reduce the brittleness sometimes found in formal games and to recognize the recurrence and interaction of broad game types (Sallach, 2006). The two aspects of interest are *pro*totypical games and *mul*tigames. Combining the two types of games results in a *promulgame*,[5] which in turn serves as an interactive building block of broader sociocultural structures and processes.

[5] Please pardon the compound neologism.

Prototype theory defines the nature of games that actors play, potentially at diverse levels. There are three prototypical games: beneficent, instrumental,[6] and coercive, with gradations of each type. Together, they define a broad range of social game types and allow exploration of the interactions among them. These games can be easily represented by a unit interval of [1, –1] divided into three ranges of, for example, [0.95 to 0.35, 0.3 to –0.3, and –0.35 to –0.95]. In this example, the games range from the most beneficent to the most coercive, with each game having an internal range of 0.6.

Promulgames are prototypical in the sense that a conceptual core defines the game type, while empirical games vary in their proximity to that core (Rosch, 1978; 1983). Researchers can identify and dimensionalize these differences for modeling purposes. This form accommodates itself to the representation of an extensive variety of empirical games, including the emergence of calibrated strategies.

Beneficent games involve a kind of mutual support seen in families and tribes, among neighbors, and within communities. Types of support in beneficent games may vary, but accounting is not strict and the game tends to be mutually reinforcing over time, commonly resulting in virtuous spirals (Carse, 1986). Instrumental games are familiar, involving complementary benefit and, relative advantage, arm's-length accounting, and are often self-reinforcing (Osborne & Rubinstein, 1990). Coercive games involve the exchange (or threat) of force or violence. Reciprocity is frequently anticipatory, and comparative accounting tends to be exaggerated, resulting in a vicious spiral. Complementarity usually takes the form of innovative tactics or novel defenses. Any of the three types of games can stabilize over time to a recurrent balance.

In addition to reciprocal interchanges, where the details primarily concern frequency, quantity, and/or quantitative conversion, off-diagonal interchanges occur in which an actor uses one resource to acquire or respond to a different one. For example, actors use force to gain material benefits, or money to gain altruistic benefits. Therefore, each game resource requires two (directional) dimensions. Such games are apparently less likely to stabilize than are reciprocal games, but are nonetheless interesting by virtue of their empirical familiarity. For example, armed robbery uses coercive means to acquire instrumental resources. Mercenary armies and/or hit men use instrumental means to acquire coercive resources. Revolutionaries and terrorists use instrumental and/or support services—providing food, welfare, and health services—in return for assistance in conducting coercive acts or hiding the perpetrators. Finally, abducting children to raise them (as opposed to demanding ransom) can be viewed as a use of coercive means of acquiring human resources.

Off-diagonal games also figure in situations characterized by ambiguity or deception as to which games are being played. Examples abound. Contributions to charities (support) may be deceptively converted to terrorist use (coercion). The sale of contraband can be an economic means of achieving the same result. Therefore, human actors (and, thus, computational models), need the

[6] This game type is often called 'economic,' but this characterization is too narrow. Other forms of instrumental and/or pragmatic trade-offs and negotiations, particularly in political and military domains, provide additional examples of instrumental games.

capability to determine the game they are being invited (enticed) to play. Computational models must assemble and parameterize instantiated games endogenously from their components (Sallach, 2000).

Social actors engage in games as a means of eliciting coordination from other actors (cf. Taylor, 2005). They initiate (or respond with) games of different types based on various situated factors. Table 1 illustrates two such factors: a) the strategy employed by the actor (top row), and b) the relationship between the actor and relevant other actors (bottom row).

Table 1. *Promulgames by Strategy (Top) and Relationship (Bottom)*

STRATEGY:	**Compliant (Soft)**	**Defensive (Firm)**	**Aggressive (Hard)**
Beneficent	Unrestrained support; give freely without expectation of return	Attend to timing and adequacy of response; be prepared to diminish reciprocity	Withdraw cooperation; shun
Instrumental	Relax competition; share opportunities; tolerate slow or incomplete responses	Compete; expect value-for-value; negotiate binding contracts	Demand hard bargains; seek to drive competitor out of region or specialty
Coercive	Show respect for adversary; accept truces; allow conflict to subside	Return blow-for-blow; assess threats; prepare for possible attacks	Attack without provocation; amplify threats; invent reasons to confront or destroy enemy
ROLE:	**Affiliate**	**Rival**	**Adversary**

The relationships among actors described in Table 1 overlap with a number of social theories and models. Recognition theory provides an important example. Honneth (1996) distinguishes among three patterns of recognition: love, rights, and solidarity. The affinity relation that Honneth calls ‘love’ gives rise to the strategies summarized in the top row, calibrated by a selected response to the particular situation. The legal recognition relationship (‘rights’) is primarily instrumental in nature (summarized in the middle row), where various actors agree to support legal rights for all as a means of securing their own (1996).

Groups based on solidarity (and its diverse bases) may draw upon any of the protogame strategies, depending on their historic and current relationships. Groups can view each other as allies (with whom they have an affinity), competitors and/or actors with asynchronous dependencies (cf.

Emerson, 1962) (with whom they engage in a pragmatic or instrumental way), or enemies (whom they view as having to be coerced, or be threatened with coercion). In each case, the calibration of particular moves depends upon the specific situation, and on the actors' understanding of it.

3. Modeling Gaps

Research has made considerable progress in modeling a wide variety of social systems, ranging from individual cognition to organizational competition. Nevertheless, these models are not perfect. Below we note several gaps for each of the models introduced in this chapter.

Blau space models offer data-efficient approaches to modeling abstract dynamics at a high level, but applying them involves more art than science. No rigorous method exists for determining which variables to select as Blau dimensions. The only theoretical guideline is that the dimensions be "salient" to the population under study, meaning that the variables must have an impact on association (Blau, 1977). In general, any variable on which association is homophilous is a salient dimension, but this is a weak criterion. Further, the modeling approach has no guidelines for the number of dimensions to include. Larger numbers of dimensions tend to reduce the overlap between niches while smaller numbers increase it, but no method exists for determining an appropriate number. This approach would therefore benefit from methods to rapidly and reliably assess whether a given Blau space model has the correct dimensions.

The types of variables that can be used as dimensions also depend on whether the analysis focuses on individuals or organizations. When depicting organizations, the models can include categorical variables (e.g., race, sex) by rendering a dimension as the percentage of a group that falls into a particular category (e.g., percentage of the organization that is female). However, this approach does not work when modeling individuals because they have discrete values (e.g., male *or* female); as a result, a Blau space model of individuals can use only continuous variables (e.g., income, age) as dimensions. New methods of scaling Blau space dimensions are needed to eliminate this restriction.

Social network models have made tremendous strides in recent decades, but require substantial improvement. First, ERGMs parameterize specific network configurations, but highly preferred configurations sometimes lack a clear meaning. In other words, a particular configuration may be necessary to achieve acceptable model fit, but there may be no clear process through which this configuration could emerge. Research should determine whether these odd configurations signal a deep insight about social behavior that remains to be achieved or simply indicate that the model is degenerate.

Additionally, while ERGMs and Siena models are powerful techniques, they both rely on simulation. This requires a considerable amount of computation time, which can be prohibitive in large and complex networks. Developers should devote additional work to optimizing these techniques for rapid computation, preferably by exploiting multi-core and parallel computing architectures whenever possible.

SINT delivers a powerful and flexible model of attitude change and stability, but several hurdles remain. First, can SINT be adapted also to model behaviors? Individuals may be more likely to adopt a particular behavior or practice as their attitude toward it improves, but no link function that relates attitudes to behaviors currently exists. It should be possible to specify such a link, but doing so would impose assumptions about the causal linkage between attitudes and behaviors. Given the pronounced issues of reverse causation in this area (i.e., attitudes may follow from behaviors), developing such a link function is a theoretical and empirical task as much as a mathematical one. Second, including an actor's attitude at the first time point in every iteration prevents perfect consensus from emerging in all cases, but also gives one arbitrary moment a disproportionate influence over the model. SINT would benefit from development of an alternative approach to preventing perfect consensus from emerging in all cases.

The majority of the models reviewed in this chapter were originally developed with face-to-face interactions in mind, but interactions increasingly occur in electronically mediated forms. It would be useful to take advantage of the wealth of online data on social behavior, but models developed around offline interactions may not translate to online contexts. Some models (e.g., SINT) have enough flexibility to handle the transition with minimal modification, but other models are likely to be context dependent. This necessarily means that we must exercise restraint in using online data until our understanding has improved, but completing this basic research will produce substantial long-term advantages, as electronically mediated communication can only continue to grow in importance.

The integration of cognitive and game-theoretic models raises four essential issues: 1) How can a broad range of social theories be integrated with each other and incorporated into computational models of social dynamics? 2) How can social models best handle cross-scale interaction? 3) How can endogenous models of situated understanding best capture the rich interactions among social actors and still remain computationally tractable? 4) What are the mathematical foundations for effective models of social dynamics? These questions define further gaps for research to address.

The first and last questions are fundamental. Answering them will provide the ultimate basis for the advancement of computational models. The second and third questions are more application oriented. They provide a focus to the more abstract questions, as well as the criteria for prospective answers. Research addressing all of these questions will result in a new generation of computational architecture, and a new foundation for computational social science.

3.1. The Road to Transition

What steps must developers take to convert the models discussed here into operational systems ready for use by practitioners? One of the most extensive, and necessary, improvements involves creating robust systems for converting between data formats. The models described above are implemented in a number of different programming languages, and one of the greatest routine difficulties in using them stems from the need to transform data to be compatible with each. If and when translation protocols emerge, developers must exercise care in minimizing the computational time required to store and transform the data.

Each of the described models has achieved significant empirical success, but, at present, guidelines for identifying degenerate or poorly fitting models remain loose. For practitioners to apply these models effectively, they must be able to determine clearly whether the models are fitting adequately and to diagnose the inevitable problems that arise in any analytic task. Therefore, developers must also create and automate diagnostics for the models described in this chapter.

Despite their separate developmental histories, many of the models presented in this chapter have deep similarities. For example, ERGMs work on networks assumed to be in equilibrium, whereas Siena models focus on disequilibrium. Similarly, both SINT and Siena focus primarily on the transition to stability, although SINT includes explicit equations for equilibrium cases. These connections suggest the possibility of integrating these models into a general framework useful in a variety of contexts. Successful integration would reduce the training time necessary for practitioners to perform many of these types of analyses competently.

This discussion of gaps defines a framework for progress in cognitive modeling and game theory. Beyond that framework, the transitional path concerns the type of computational architecture in which these questions cannot only be addressed, but can also evolve as innovations give rise to new issues. The broad computational form will probably be based on agent-based modeling and simulation, a topic deemed beyond the scope of the present chapter. However, developers should view the present forms of agent simulation as bare bones relative to the capabilities ultimately required. The framework defining the gaps also identifies the transitional path. We can expect that agent simulation in the future will be far more theoretical and extensively grounded in mathematical relations and transformations (cf. Sallach, 2012). As game formality increases, it will not only specify simulations, but also form a basis for deductive analysis. Each of these advances will transform and extend computational social science, allowing it to realize its full potential.

4. Conclusions

Research has made substantial progress in developing and refining computational models of human behavior, and has produced a palette of models useful for a broad range of specific contexts and data limitations. Additional work is needed—both basic research to improve our understanding of the workings and implications of the models and applied research to make them easier for practitioners to use—but the basic groundwork has been completed.

The models discussed in this chapter aid in understanding the operation of social groups in a variety or respects. They can provide clues as to the processes that yield specific network configurations and the likelihood of different configurations, can describe the degree of competition between organizations and other abstract entities, and can explain the inner workings of attitude change. We must reemphasize that each model described here is useful in particular circumstances and depends on specific assumptions. Practitioners should therefore assess the usefulness of a model on the basis of these assumptions and of their validity in a particular modeling context. Assumptions will often be strained, if not actually violated, but practitioners should take care to ensure that these violations remain within acceptable bounds. This chapter has surveyed many models, and made no attempt to describe the various assumptions made exhaustively, but interested readers can consult the literature cited to learn more about these specific techniques and their advantages and disadvantages.

Computational models lend themselves to diverse theories and their integration. Even as researchers assess the theories that animate them more broadly and theoretically, the models must retain their expressiveness. Drawing on the insights attained, their formulation as mathematical models, more elegant and more compelling, will contribute greatly to the advance of mathematical sociology. In 20 years, we can expect that social models will have become considerably more effective because of the mutual strengthening of theoretical sociology and the mathematical formulations of these theories.

References

Becker, G. (1976). *The economic approach to human behavior.* Chicago, IL: University of Chicago Press

Behler, R., Suh, C. S., & Brashears, M. E. (2013). *The influence of indirect others: An ecological approach to understanding weight-related perceptions and behaviors*. Unpublished manuscript.

Benmelech, E., & Berrebi, C. (2007). Human capital and the productivity of suicide bombers. *The Journal of Economic Perspectives, 21*(Summer): 223-238.

Bewley, T. F. (1999). *Why wages don't fall during a recession.* Cambridge, MA: Harvard University Press.

Blau, P. M. (1977). *Inequality and heterogeneity: A primitive theory of social structure*. New York, NY: The Free Press.

Bonikowski, B. (2010). *Ecology of taste: A dynamic analysis of musical genre preferences, 1982-2002.* Unpublished manuscript.

Brandes, U., Indlekofer, N., & Mader, M. (2012). Visualization methods for longitudinal social networks and stochastic actor-oriented modeling. *Social Networks*, *34*, 291-308.

Brashears, M. E. (2008). Sex, society and association: A cross-national examination of status construction theory. *Social Psychology Quarterly*, *71*, 72-85.

Brashears, M. E., Genkin, M. & Suh, C. S. (2012). *In the organization's shadow: How individual behavior is shaped by organizational competition.* Unpublished Manuscript.

Bratman, M. E. (1987). *Intention, plans, and practical reason.* Stanford, CA: CSLI Publications.

Caimo, A. & Friel, N. (2013). Bayesian model selection for exponential random graph models. *Social Network,* 35, 11–24.

Carse, J. P. (1986). *Finite and infinite games*. New York, NY: Ballantine.

Coleman, J. (1990). *Foundations of social theory.* Cambridge, MA: Harvard University Press.

Collins, R. (1981a). On the microfoundations of macrosociology. *American Journal of Sociology, 86*, 984–1014.

Collins, R. (1981b). *Sociology since midcentury: Essays in theory cumulation*. New York, NY: Academic Press.

Collins, R. (1989). Sociology: Proscience or antiscience? *American Sociological Review, 54*, 124–139.

Collins, R. (1990). Stratification, emotional energy and the transient emotions. In T. D. Kemper (Ed.), *Research agendas in the sociology of emotions* (pp. 27–57). Albany, NY: SUNY Press.

Collins, R. (1993). Emotional energy as the common denominator of rational choice. *Rationality and Society, 5*, 203–230.

Collins, R. (2005). *Interaction ritual chains.* Princeton, NJ: Princeton University Press.

Doran, J. E. (2000). Trajectories to complexity in artificial societies: Rationality, belief, and emotions. In T.A. Kohler & G.J. Gumerman (Eds.), *Dynamics in human and primate societies: Agent-based modeling of social and spatial processes* (pp. 89-105). New York, NY: Oxford University Press.

Durkheim, E. (1995). *The elementary forms of religious life.* New York, NY: Free Press.

Elster, J. (Ed.). (1986). *Rational choice.* New York, NY: New York University Press.

Field, A. J. (1984). Microeconomics, norms and rationality. *Economic Development and Cultural Change, 32*, 683-711.

Fisher, F. M. (1989). Games economists play: A non-cooperative view. *RAND Journal of Economics, 20*, 113-124.

Fudenberg, D., & Tirole, J. (1987). Understanding rent dissipation: On the use of game theory in industrial organization. *American Economic Review, 77*, 176-183.

Friedkin, N. E., & Johnsen, E. C. (2011). *Social Influence network theory: A sociological examination of small group dynamics*. Cambridge, MA: Cambridge University Press.

Granovetter, M. S. (1995). *Getting a job: A study in contacts and careers.* Chicago, IL: University of Chicago Press.

Handcock, M. S. (2003). Statistical models for social networks: Degeneracy and inference. In Breiger, R., Carley, K., & Pattison, P. (Eds.), *Dynamic Social Network Modeling and Analysis,* (pp. 229-240). Washington, D.C.: National Academies Press.

Honneth, A. (1996). The struggle for recognition: The moral grammar of social conflicts. Cambridge, MA: MIT Press.

Kemper, T. D., & Collins, R. (1990). Dimensions of microinteraction. *American Journal of Sociology, 96*, 32-68.

Lazarsfeld, P. F. & Merton, R. K. (1954). Friendship as a social process: A substantive and methodological analysis. *Freedom and Control in a Modern Society, 18*, 18-66.

Levy, J. S. (1997). Prospect theory, rational choice and international relations. *International Studies Quarterly 41*, 87-112.

Lusher, D., Koskinen, J., & Robins, G. (2013). *Exponential random graph models for social networks: Theory, methods, and applications*. New York, NY: Cambridge University Press.

Luttwak, E. N. (1987). *Strategy: The logic of war and peace*. Cambridge, MA: Harvard University Press.

Mark, N. (1998). Birds of a feather sing together. *Social Forces, 77*, 453-485.

Mark, N. (2003). Culture and competition: Homophily and distancing explanations for cultural niches. *American Sociological Review*, *68*, 319-345.

Marsden, P. V. (1987). Core discussion networks of americans. *American Sociological Review, 52,* 122-131.

Marsden, P. V. (1988). Homogeneity in confiding relations. *Social Networks*, *10*, 57-76.

Marsden, P. V. & Friedkin, N. E. (1993). Network studies of social influence. *Sociological Methods & Research*, *22*, 127-151.

Marsden, P. V., & Gorman, E. H. (2001). Social networks, job changes and recruitment. In I.E. Berg & A.L. Kallenberg (Eds.), *Sourcebook of labor markets: Evolving structures and processes*. (pp. 503-530). New York, NY: Kluwer Academic/Plenum Publishers.

McPherson, J. M. (1983). An ecology of affiliation. *American Sociological Review 48*:519-532.

McPherson, J. M. (2004). A Blau space primer: Prolegomenon to an ecology of affiliation. *Industrial and Corporate Change*, *13*, 263-280.

McPherson, J. M., & Ranger-Moore, J. R. (1991). Evolution on a dancing landscape: Organizations and networks in dynamic Blau space. *Social Forces*, *70*, 19-42.

McPherson, J. M., & Rotolo, T. (1996). Testing a dynamic model of social composition: Diversity and change in voluntary groups. *American Sociological Review*, *61*, 179-202.

McPherson, J. M., Smith-Lovin, L., & Cook, J. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*, 415-444.

Miller, G. F. (1997). Protean primates: Adaptive unpredictability in courtship and competition. In A. Whiten & R.W. Byrne, (Eds.). *Machiavellian intelligence II: Extensions and evaluations*. (pp. 312-340). New York, NY: Cambridge University Press.

Osborne, M. J. & Rubinstein, A. (1990). *Bargaining and markets*. San Diego, CA: Academic Press.

Pape, R. (2005). Dying to win: The strategic logic of suicide terrorism. New York, NY: Random House.

Parsons, S., Gmytrasiewicz, P., & Wooldridge, M. (Eds.). (2002). *Game theory and decision theory in agent-based systems.* Boston, MA: Kluwer Academic.

Peltzman, S. (1991). The handbook of industrial organization. *Journal of Political Economy, 99,* 201-217.

Popielarz, P. A. & McPherson, J. M. (1995). On the edge or in between: Niche position, niche overlap and the duration of voluntary association memberships. *American Journal of Sociology*, *101*, 698-720.

Rao, A. S., & Georgeff, M. P. (1995, June). BDI agents: From theory to practice. In *Proceedings of the First International Conference on Multi-Agent Systems.* San Francisco, CA.

Robins, G., Pattison, P., Kalish, Y., & Lusher, D. (2007). An introduction to exponential random graph (p^*) models for social networks. *Social Networks, 29*, 173-191.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B.B. Lloyd (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Lawrence Erlbaum.

Rosch, E. (1983). Prototype classification and logical classification. In E.K. Scholnick (Ed.), *New trends in conceptual representation: Challenges to Piaget's theory?* (*73-86).* Hillsdale, NJ: Lawrence Erlbaum.

Sallach, D. L. (2000, June). Games social agents play: A complex form. In *Joint Conference on Mathematical Sociology in Japan and America*. Honolulu, HI.

Sallach, D. L. (2006). Complex multigames: Toward an ecology of information artifacts. In D. L. Sallach, C. M. Macal, & M. J. North (Eds.), *Proceedings from the Agent 2006 Conference on Social Agents: Results and Prospects*. Chicago, IL: Argonne National Laboratory.

Sallach, D. L. (2008). Modeling emotional dynamics: currency versus field. *Rationality and Society, 20*, 343-365.

Sallach, D. L. (2012, September). Categorical social science: Theory, methodology and design. *Proceedings of the World Congress on Social Simulation.* Taipei, Taiwan.

Snijders, T. A. B., Pattison, P. E., Robins, G. L. & Handcock, M. S. (2006). New specifications for exponential random graph models. *Sociological Methodology, 36*, 99-153.

Snijders, T. A. B., van de Bunt, G. G. & Steglich, C. E. G. (2010). Introduction to stochastic actor-based models for network dynamics. *Social Networks*, *32*, 44-60.

Snijders, T. A. B., Lomi, A., & Torlo, V. (2013). A model for the multiplex dynamics of two-mode and one-mode networks, with an application to employment preference, friendship, and advice. *Social Networks, 35*, 2, 265-276.

Suh, C. S., Brashears, M. E. & Genkin, M. (2012). *Who are the party people? Organizational membership and adolescents' use of recreational drugs*. Unpublished manuscript.

Taylor, A. D. (2005). *Social choice and the mathematics of manipulation.* New York, NY: Cambridge University Press.

Wang, P., Robins, G., Pattison, P. & Lazega, E. (2013). Exponential random graph models for multilevel networks. *Social Networks, 35*(1), 95-116.

Wasserman, S. & Faust, K. (1999). *Social network analysis: Methods and applications.* Cambridge, MA: Cambridge University Press.

Wasserman, S. & Robins, G. L. (2005). An Introduction to random graphs, dependence graphs, and p^*. In P. Carrington, J. Scott, & S. Wasserman (Eds.), *Models and methods in social network analysis (*148-161). New York, NY: Cambridge University Press.