In Proceedings of the Sixth International Conference on Cognitive Modeling, 290-295.
Mahwah, NJ: Lawrence Earlbaum.

A Dynamic, Multi-Agent Model of Peer Group Formation

Terrence C. Stewart (tcstewar@connect.carleton.ca)
Institute of Cognitive Science, Carleton University
1125 Colonel By Drive, Ottawa, ON, K1S 5B6, Canada

Robert L. West (robert_west@carleton.ca)
Institute of Cognitive Science; Department of Psychology, Carleton University
1125 Colonel By Drive, Ottawa, ON, K1S 5B6, Canada

Robert Coplan (robert_coplan@carleton.ca)
Department of Psychology, Carleton University
1125 Colonel By Drive, Ottawa, ON, K1S 5B6, Canada

Abstract

A computational model of the formation of peer groups in children is presented. We used standard sociometric measurements (the CDC classification of Popular, Rejected, Neglected, and Controversial individuals) to compare the model to empirical data. The model fit this data well in terms of category distributions and stability, even without introducing individual differences. When individual differences were added, the model went on to accurately predict a separate set of empirical results. Furthermore, patterns arose in the results which suggested an underlying relationship between certain types of individual differences.

Introduction

In recent years, we have started to see an increasing willingness among researchers to examine high-level models of human group behavior. Interesting results have been found regarding wealth distribution (Bouchaud & Mézard, 2000), human rioting (Granovetter, 1978), crowd panic (Helbing, Farkas, & Vicsek, 2000), and friendship networks (Watts & Strogatz, 1998), to name just a few. We have also seen the development of larger frameworks for investigating social cooperation to achieve goals in general (for example, Sun, 2001). These models provide new insight into these fields, along with new predictions that can be tested.

In this paper, we present a computational model of the formation of friendship groups among peers. In the model, individuals meet and interact over time, becoming friends (and possibly enemies) with each other. This process has been empirically investigated by developmental psychologists, who have developed standardized measures of popularity (see Cillessen & Bukowski, 2000, for a review). In this paper we will evaluate our model's results by comparison to the CDC measure (named after its creators' initials: Coie, Dodge, & Coppotelli, 1982), which is the most commonly used system.

Psychologists do not consider 'popularity' to be an intrinsic aspect of an individual. Rather, 'popularity' is a dynamic construct formed among a group of individuals. That is, an individual's popularity is determined by how others feel about them. However, explanations as to why a person is popular or unpopular focus on individual characteristics, particularly personality factors. Individual factors do influence popularity; a child could be unpopular because they are shy, aggressive, and/or socially incompetent (Newcomb, Bukowski, & Pattee, 1993). However, Rubin, Hymel, & Mills (1989) reviewed the findings in this area and found that only a small percentage of the variance could be accounted for by individual differences.

While it is possible that the unaccounted for variance is simply due to an inability to measure popularity and/or personality factors well enough, another possibility is that it can be attributed to the dynamics of the system. That is, people's popularity may have more to do with the dynamics of their peer group interactions than with their own personality. If so, this would mark a radical departure from the common understanding of popularity as a product of personality. We investigated this possibility through the use of multi-agent modeling.

The CDC Measure

The goal of the CDC measure is to classify people into one of five categories: *Popular*, *Rejected*, *Controversial*, *Neglected*, or *Average*. Using interviews or questionnaires, each person is asked to name three people in their peer group that they like, and three people that they dislike. The simplicity of this measurement is important for measuring popularity in young age groups. Using the survey results, each individual is given an *Acceptance* score (the total number of times that person is listed by other people as someone they like) and a *Rejection* score (the number of times they appear on the 'dislike' lists). A *Preference* value (Acceptance minus Rejection) and an *Impact* value (Acceptance plus Rejection) are also created; where preference refers to whether you are more liked or disliked and impact refers to how much people pay attention to you. Individuals are then classified into the five categories according to the rules shown in Table 1.

Table 1: The decision rules for classifying with CDC. All values are normalized to a mean of 0 and a standard deviation of 1.

Category	Rule
Popular	Preference>1 Acceptance>0 Rejection<0
Rejected	Preference<-1 Acceptance<0 Rejection>0
Neglected	Impact<-1 Never appears on anyone's 'like' list
Controversial	Impact>1 Acceptance>0 Rejection>0
Average	None of the Above

Comparison Data

Newcomb, Bukowski, & Pattee (1993) provide a considerable amount of data to test our model against. To begin with, they give complete results from nine different studies that used CDC categorization on a total of 2,571 students, ranging from kindergarten to grade 9. These nine studies give the following 95% confidence intervals for the percentage of people in each category.

Table 2: The 95% confidence intervals indicating the percentage of people in each CDC category.

Popular	Rejected	Neglected	Contro.	Average
7%-32%	12%-26%	0-28%	1.6-16%	5.9-69%

This data set gives us our first basis of comparison between the model and reality. In particular, if we find that a model gives results outside of this range, then we can conclude that it is not a suitable model.

The next set of data that we can use for comparison is the stability of these measurements over time. Cillessen, Bukowski, & Haselager (2000) give the results of a meta-study which collected the change in CDC categorization over periods of time ranging from one month to four years. The results give the percentage of people who were in the same category at the beginning and at the end of the study period. The following figure gives the 95% Confidence Interval for this data.

Table 3: The 95% confidence intervals for the stability of each CDC category.

Popular	Rejected	Neglected	Contro.	Average
33-44%	39-49%	20%-30%	24%-36%	51%-69%

In addition, the results showed no significant change related to the amount of time in the study period. In other words, there is no significant difference between the stability over four months and stability over four years. We should thus expect a good computational model to not only match the stability data shown in table 3, but also to give this result independent of the number of interactions simulated.

Creating a Benchmark

To further evaluate our model, we created a benchmark for evaluating CDC results. Specifically, we wanted to create a null condition result, representing "no effect" for the CDC measure. To do this used a completely random response scheme. That is, we determined what would happen if the 'like most' and 'like least' responses were generated by having each person nominate three others completely at random. This data was then subjected to the standard CDC analysis, as described. We chose a group size of 30 to be consistent with real-life situations.

After generating 1000 groups, we collected the CDC classification data and found the following distribution.

Table 4: The distribution of individuals in the benchmark.

Popular	Rejected	Neglected	Contro.	Average
12%	12%	7%	2%	67%

Surprisingly, the null condition results did not fall outside the confidence intervals for the categorization data in table 2. Thus the categorization distribution results cannot be considered to be different from chance. When we turn to the stability data, however, we do find a statistically significant difference. Given the simplicity of this random model, since an individual has a 12% chance of being classified as Popular, when we re-run the simulation that same person will still have only a 12% chance of being Popular this time. This means that the stability data will be identical to the previous table

Table 5: The stability of categorization in the benchmark.

Popular	Rejected	Neglected	Contro.	Average
12%	12%	7%	2%	67%

These results were well outside the confidence intervals presented in table 3 ($p<0.001$). The null condition clearly does not capture important aspects of the process. The results for the stability data show that the CDC findings are meaningful. However, the results for the categorization data show that the apparent differences in the sizes of the categories can be viewed as an artifact of the way the CDC system works.

Given this baseline, we can see that predicting the category distribution data (Table 2) would only provide weak support for a model. However, if a model was able to predict both the distribution and the stability data, then we could be confident that it is matching the real-world data well.

Our Model

In developing our model, we tried to determine the simplest system possible that had the following characteristics:

1. Each person should remember how much they like each other person

2. People should use this memory to determine how they will 'interact' with others
3. People should use the results of this 'interaction' to change how much they like the person they just interacted with

This led us to the following algorithm:

1. Let a[i,j] be the amount person i likes person j
2. For all pairs of individuals i,j:
   a. Have i use a[i,j] to decide how to behave towards j
   b. Have j use a[j,i] to decide how to behave towards i
   c. Update a[i,j] based on how j behaves towards i
   d. Update a[j,i] based on how i behaves towards j
3. Repeat (2) for as many iterations as desired

To implement this, we needed to define an algorithm for each agent to use to decide how to behave towards another, based on how much they like each other (steps 2a and 2b). We chose a simple method: take the value of a[i,j], add a random value from a Gaussian distribution, and use the resulting value to represent how 'good' individual i is going to behave towards j.

$$b[i,j]=a[i,j]+N(0,1) \qquad \text{(formula 1)}$$

The random variable has a deviation of 1. Changing this deviation does not affect the behavior of the model.

Similarly, we needed to update how much i likes j, based on j's actions. One idea would be to simply add b[j,i] (how 'nicely' j has behaved toward i) to a[i,j]. However, this approach can be easily shown to cause a positive feedback loop which would mean that if i currently likes j and j currently likes i, they will keep increasing how much they like each other to astronomical values. This is clearly does not capture the ebb and flow of real human relations. Instead, we use a slightly more complex formula which acknowledges the role that a[i,j] (the amount person i likes person j) could play in determining how person i would react to the behavior of person j.

$$\Delta a[i,j] = w_{lr}(b[j,i]-a[i,j]) \qquad \text{(formula 2)}$$

The result of this was that individuals in the model evaluated the actions of others according to the expectations they had for those others. This was based on the insight that we expect a friend to be friendly and an enemy to be unfriendly, and we are not surprised when this happens. However, if a friend were to be unfriendly we could be quite hurt. Likewise, if an enemy were to be kind it would surprise us and could cause us to reevaluate our feelings toward them.

Since repeated iterations of this model will cause a[i,j] and a[j,i] to interact with one another, it is instructive to view what typically happens to these variables over time. Figure 1 illustrates the effect. The results show that the model can potentially capture the ups and downs of real human relationships.

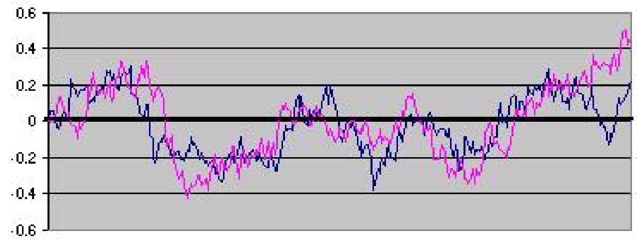

Figure 1: Two individuals interacting over 300 iterations of the model. Shown is a[i,j] and a[j,i]. w_{lr} is 0.1.

It should be noted that in doing this we have introduced the first parameter into this model: w_{lr}. This is the 'learning rate' and controls how quickly an individual changes its opinions of others. It can be considered to be similar to the learning rate parameters found in a wide variety of other computational models. This is a very interesting parameter because creates a potential link between cognitive abilities and popularity, something that has not been previously explored. Essentially, the model examines the role of memory, learning and expectations in human relationships.

Category Distribution Results

Interestingly, it turns out to be very easy to show that this new model is exactly as good as the random model at predicting the distributions of categories. We first note that, given the preceding algorithm, the values a[i,j] and a[j,i] change in a manner that is independent of anything else in the simulation. This means that when the CDC nominations occur (i.e. when we check each individual to see which three others they like and dislike the most), the chance of a particular person j nominating another person i is independent of person k also nominating person i. This in turn means that, for any given CDC evaluation, the number of nominations an individual receives (both positive and negative) will have the same sort of distribution as the initial completely random model. This predicts that we will continue to have exactly the same distributions of categories in this new model as under the random model. We ran the simulations anyway, and found that we did, in fact, get the same distributions (for all values of w_{lr}).

Table 5: The distribution of individuals in the model.

Popular	Rejected	Neglected	Contro.	Average
12%	12%	7%	2%	67%

Category Stability Results

We can now determine how well this model predicts the category stability data. In order to do this, we first need to choose a value for w_{lr} (the learning rate which controls the speed of adaptation of an individual's 'liking' of another). Possible values range from 0 to 1. To deal with this, we can run the simulation for multiple values of w_{lr}.

We also need to determine how many iterations our model should be run for. Given the real-world data from Cillessen, Bukowski, & Haselager (2000), we know that a good model will result in asymptotic stabilities. That is, over time the stability of the Popular category should approach a value between 33% and 44%. Similar expectations exist for the other four categories.

Figure 2 shows our model's stability results for each of the five CDC categories. Each line in each graph represents a different setting for w_{lr}. The x-axis indicates the number of iterations (i.e. time) between the two category measurements. The far right gives our results after 20 time steps. Overlaid on each of these graphs is the 95% confidence interval from the real-life stability data.

Analyzing this data, we can see that the precise value of w_{lr} does not significantly affect our results. That is, we can set our learning rate parameter anywhere within its possible range and still get very similar results. Our model is thus effectively a zero-parameter model.

The other promising result is that each of the five stability measurements end up in or near the 95% Confidence Interval. This is a very encouraging finding, since no parameter tweaking was required.

Individual differences

Although our model matches the data well without introducing individual differences, it is still the case that these differences have been found to have some influence on popularity. To investigate this factor, we systematically introduced different sorts of variations in our model to investigate four well-known effects. In all of these simulations, we allowed a particular aspect of each individual in the model to vary across a normal distribution. The size of this distribution was made as large as possible while still matching to the real-life data.

To begin with, we looked at the Hostile Attribution Bias. This refers to the fact that certain individuals tend to interpret the actions of others in a more negative light than is intended. Research has shown that rejected children (particularly those who are aggressive) are more likely to assume malevolent intent when they are faced with ambiguous social cues (Crick & Dodge, 1994; Dodge, Lansford, Burks, Bates, Pettit, Fontaine, & Price, 2003). To model this we inserted this individual difference into our model by adjusting formula (2) in the following way:

$$\Delta a[i,j] = w_{lr}(b[j,i] + B[i] - a[i,j])$$

In this new formula, B is an individual interpretation bias. Note that in our simulation, agents could be biased to be either overly negative or overly positive in how they interpreted the actions of others.

We then ran the simulation over 1000 groups of 30 agents each. Each agent had a value of B chosen from a normal distribution with a deviation as large as possible while still matching the aforementioned results. After 50 simulation iterations, CDC classification was performed. We then measured the effect size. This involved determining how many standard deviations above or below the mean the values of B were within each category. This method was used for all individual differences investigated.

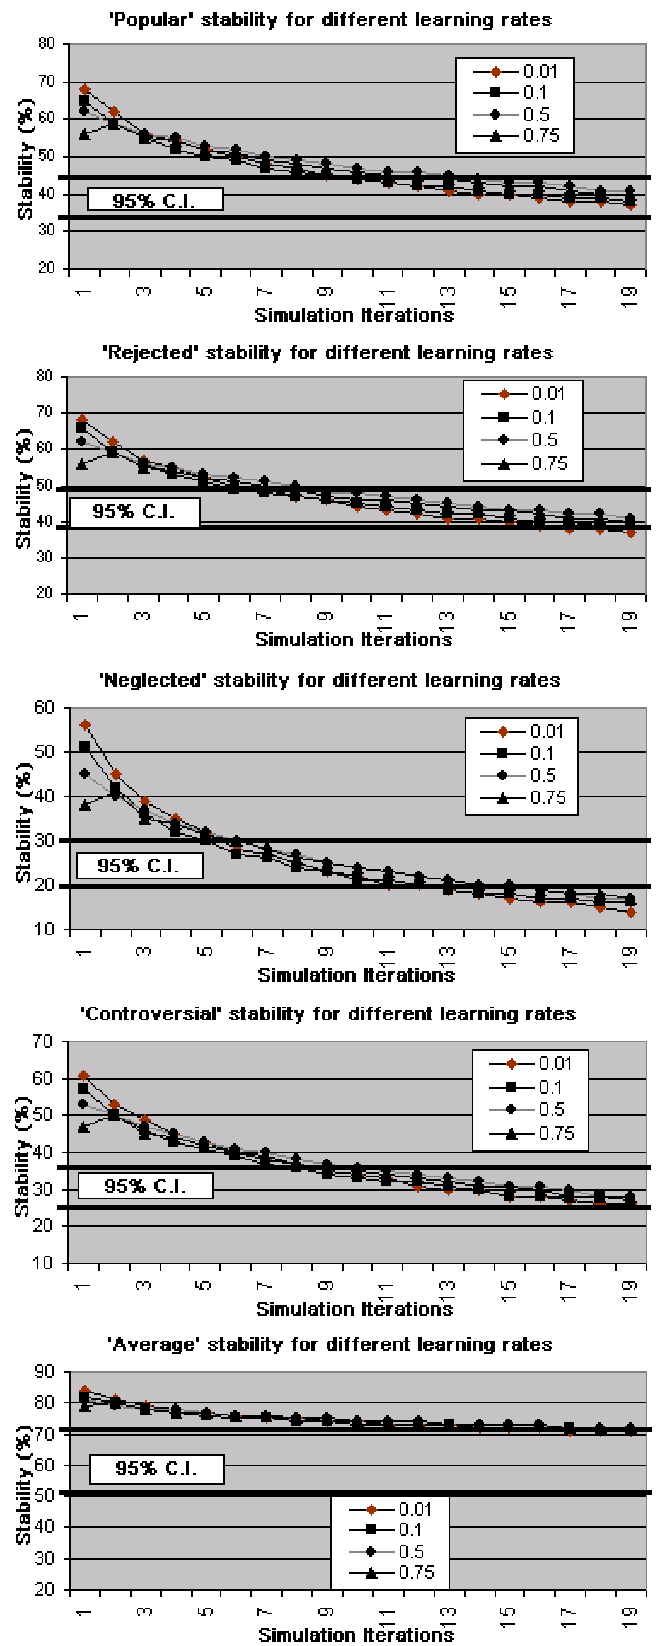

Figure 2: Simulation stability results compared to real-world data from (Cillessen, Bukowski, & Haselager, 2000), compared over a range of settings for w_{lr}.

The results in Figure 3 show that, as predicted, the Rejected agents tended to be overly negative in interpreting the behaviors of others, while the Popular agents tended to be overly positive. It was also interesting that the Neglected agents also tended to be somewhat negative.

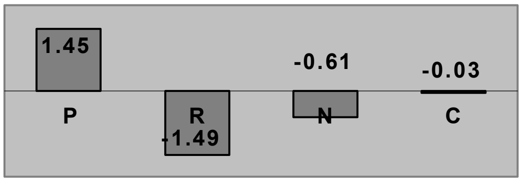

Figure 3: Effect size of varying interpretation bias (the average number of deviations a group is from the mean)

Another well-known finding is that Popular children tend to have good social skills whereas rejected children tend to have poor social skills. (e.g., Coie, Dodge, & Kupersmidt, 1990; Newcomb, Bukowski, & Pattee, 1993; Parkhurst & Hopmeyer, 1998). To reflect this in the model we biased the mean of the Gaussian distribution used in Formula 1. Specifically, a positive value created a bias toward behaving nicely and a negative value created a bias toward behaving badly. Figure 4 displays the results and shows that the manipulation had the intended effect. Interestingly, the neglected agents were again shown to be somewhat similar to the rejected agents.

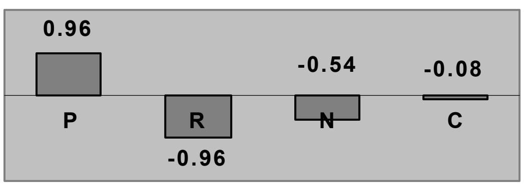

Figure 4: Effect size of varying behavior bias (the average number of deviations a group is from the mean)

The third effect we investigated was that neglected children have been shown to interact with their peers less frequently than average children (Dodge, Coie, & Brakke, 1982; Coie & Dodge, 1988). To reflect this in our model, we added an interaction probability for each agent. The percentage chance for two agents interacting was determined by multiplying their interaction probabilities together. As illustrated in Figure 5, this manipulation was successful in capturing the effect. However, it also produced the unexpected effect that a high level of interaction was associated with being controversial.

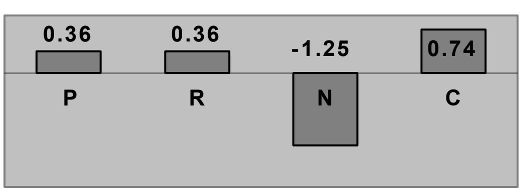

Figure 5: Effect size of varying interaction probability (the average number of deviations a group is from the mean)

Finally, we examined the finding that controversial children tend to display a combination of positive and negative social behaviors (e.g., Coie & Dodge, 1988). We modeled this lack of consistency by varying the standard deviation of the Gaussian distribution in Formula 1. The results, displayed in Figure 6 supported our interpretation that Controversials tend to be highly variable in their behavior. The results also revealed an unexpected effect in which neglected individuals tended to be more reliable in their behavior (lower variability).

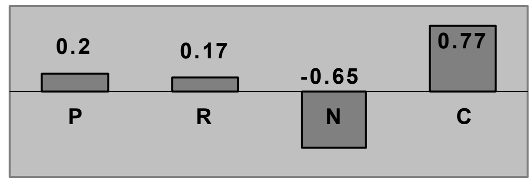

Figure 6: Effect size of varying behavior deviation (the average number of deviations a group is from the mean)

Insights and Predictions

The model also gave us the opportunity to explore other possible individual differences. The first that we looked at was the initial value for a[i,j]. For this simulation some agents started off more predisposed to liking everyone (a high value of a[i,j] for all j), and others more predisposed to disliking everyone (a low value of a[i,j] for all j). This was meant to represent the effect of previous experience, before entering the group (e.g., family experiences or other peer group experiences). The results of this simulation are displayed in Figure 7. Note that the results of this simulation were very similar to the results of the hostile attribution simulation and the social skills simulation. All of these showed an expected association between strong positive biases and being popular, and strong negative biases and being rejected. But they also all showed a somewhat unexpected association between moderate negative biases and being neglected. Taken together, these results show that the model is very robust in producing this pattern of results in response to factors that bias an agent to be more positive or more negative.

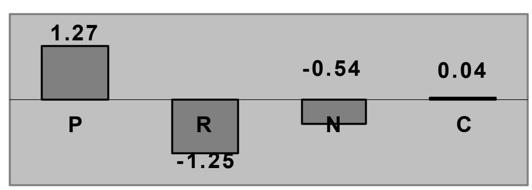

Figure 7: Effect size of varying the initial value of a[i,j] (the average number of deviations a group is from the mean)

We also analyzed the effect of individual differences on $w_{lr.}$ That is, how does having a different learning rate affect one's eventual CDC categorization. The results, displayed in Figure 8, were interesting, in that they were very similar to the results of the interaction probability simulation and the behavior consistency simulation. The most straightforward interpretation of this is that having a high learning rate or having many interactions can have the effect of making one appear to be more variable in behavior. This result is also important since we earlier (see Figure 2) showed that varying the overall learning rate among all the individuals had no effect.

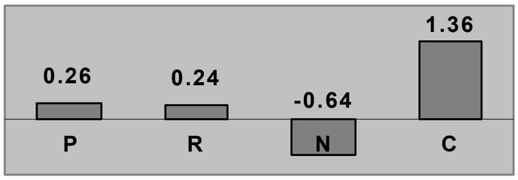

Figure 3: Classification effect size of varying w_{lr} (the average number of deviations a group is from the mean)

Conclusions

This model has presented us with a number of intriguing results. We have shown that the distribution of individuals into the five CDC categories is an artifact of the measurement system itself. The model we have presented predicts both the distribution data and the long-term stability data of these categories very well (the model predicts Neglected individuals have a stability which is 5% lower than the real-life data, but for all other categories it is within the 95% confidence interval). Our model can do this without recourse to any individual differences. Instead, these categorizations arise from the dynamics of the interpersonal interactions. These results hold over all settings for the model's one parameter, which means it is effectively a zero-parameter model.

When we introduced individual differences into the model, it predicted a number of standard effects of such differences in real children. Furthermore, these effects indicated underlying similarities in the processes involved. For example, we have consistently seen that changes which correlate with an individual being Rejected tend to also have a smaller effect associated with being Neglected. Also, we showed a striking similarity between the results of varying the learning rate, the behavioral variability, and the interaction probability. These patterns suggest new ways of looking at current research in popularity.

References

Bouchaud, J-P. & Mézard, M. (2000). Wealth condensation in a simple model of economy, *Physica A*, *282*, 536.

Cassidy, J., Scolton, K.L., Kirsh, S.J., & Parke, R.D. (1996). Attachment and representations of peer relationships. *Developmental Psychology, 32*, 892-904.

Cillessen, A., Bukowski, W., & Haselager, G. (2000). Stability of Sociometric Categories. *New Directions for Child and Adolescent Development, 88*, 75-93.

Coie, J.D. & Dodge, K.A., (1988). Multiple sources of data on social behavior and social status. *Child Development, 59*, 815-829.

Coie, J.D., Dodge, K.A., & Kupersmidt, J.B. (1990). Peer group behavior and social status. In S.R. Asher & J.D. Coie (Eds.), *Peer rejection in childhood*. New York: Cambridge University Press.

Cohn, D.A. (1990). Child-mother attachment of six-year-olds and social competence at school. *Child Development, 61*, 152-162.

Crick, N. R., & Dodge, K. A. (1994). A review and reformulation of social information-processing mechanisms in children's social adjustment. *Psychological Bulletin, 115*, 74-101.

Dodge, K.A., Coie, J.D., & Brakke, N.P. (1982). Behavioral patterns of socially rejected and neglected preadolescents: The roles of social approach and aggression. *Journal of Abnormal Child Psychology, 10*, 389-409.

Dodge, K. A., Lansford, J., Burks, V., Bates, J.E., Pettit, G., Fontaine, R., & Price J.M. (2003). Peer rejection and social information-processing factors in the development of aggressive behavior problems in children. *Child Development, 74*, 374-393

Elicker, J., Englund, M., & Sroufe, L.A. (1992). Predicting peer competence and peer relationships in childhood from early parent-child relationships. In R. Parke & G. Ladd (Eds.), *Family-peer relationships: Models of linkage*. Hillsdale, NJ: Erlbaum.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology. 83:6,* 1420-1443.

Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating Dynamical Features of Escape Panic. *Nature 407*, 487-490.

Newcomb, A.F., & Bukowski, W.M., & Pattee, L. (1993). Children's peer relations: A meta-analyic review of popular, rejected, neglected, controversial, and average sociometric status. *Psychological Bulletin, 113*, 99-128.

Parkhurst, J.T., & Hopmeyer, A. (1998). Sociometric popularity and peer-perceived popularity: Two distinct dimensions of peer status. *Journal of Early Adolescence, 18*, 125-144.

Rubin, K.H., Hymel, S., & Mills, R.S. (1989). Sociability and social withdrawal in childhood: Stability and outcomes. *Journal of Personality, 57*, 237-255.

Sun, R. (2001). Cognitive science meets multi-agent systems: a prolegomenon. *Philosophical Psychology*, *14:1*, 5-28.

Watts, D.J. & Strogatz, S.H. (1998). Collective dynamics of 'small-world' networks. *Nature, 393*, 440-442.