Alcohol Use among Adolescents as a Coordination Problem in a Dynamic Network *

Rense Corten † Andrea Knecht ‡

June 25, 2010

1 Introduction

Adolescence is a life stage in which many forms of problematic behavior reach their peak (Steinberg and Morris, 2001), among which delinquency and substance abuse. Even though there is little evidence that problematic behavior in adulthood *originates* from behavior during adolescence (Moffitt, 1993), these types of behavior may have considerable impact on adolescents themselves and on society in general. Especially substance abuse, on which we focus in this paper, is associated with problems in other areas, including delinquency, mental health problems, and problems with educational attainment (Newcomb and Bentler, 1989). Social influence by peer groups has often been named as one of the important factors that can trigger various types of problematic behavior, including alcohol and drug abuse, (e.g., Hawkins et al., 1996; Moffitt, 1993; Newcomb and Bentler, 1989). Consequently, relationships among adolescents have been the focus of a considerable body of literature (Giordano, 2003).

Issues of peer influence and selection have been of major concern within this context. On the one hand, it has been found that adolescents are sensitive to the influence of peers (Graham et al., 1991; Swadi, 1999; Bot et al., 2005). On the other hand, it has also been recognized that not only are adolescents influenced by their social environment, but also choose peers as friends who are similar to themselves, leading to network homophily (Lazarsfeld and Merton, 1954; McPherson et al., 2001). Disentangling these two simultaneous processes poses an ongoing theoretical and methodological challenge (Bauman and Ennett, 1996; Kirke, 2004).

*This paper is work in progress. We thank Vincent Buskens, Werner Raub, Stephanie Rosenkranz and Jeroen Weesie for constructive comments and suggestions.

†Corresponding author. Department of Sociology/ICS, Utrecht University. E-mail: r.corten@uu.nl.

‡Department of Sociology, University of Erlangen-Nuremberg. E-mail: andrea.knecht@wiso.uni-erlangen.de.

Small and medium-sized groups of adolescents have since long been a popular setting for sociologists to study social networks, and the emergence of cultures and norms. Schools, in particular, are an attractive setting to study these topics because they constitute relatively well-delineated social contexts, in which complex processes can be observed relatively easily (some prominent examples include Coleman, 1961; Epstein and Karweit, 1983; Bearman et al., 2004). In this sense, a school constitutes a kind of "social microcosm," or, as Coleman (1961, p. 9) put it, a "society within society."

Research on diffusion dynamics has shown that the overall structure of a network has important consequences for emerging patterns of behavior (e.g., Granovetter, 1973; Watts, 2002; Centola and Macy, 2007). Yet, while studies on adolescent behavior often emphasize the importance of social networks, most research focuses on *individual-level* explanatory factors such as attributes of personal networks (e.g., Graham et al., 1991) or individual network positions (e.g., Ennett and Bauman, 1993). A possible reason for this divergence is that theorizing on the effects of the macrolevel network structure requires the specification of the mechanisms that connect the macrolevel network structure with microlevel individual behavior, and conversely, microlevel behavior with macrolevel collective outcomes. Given the interdependencies involved in interpersonal influence processes, specifying these macro-to-micro and micro-to-macro mechanisms is not trivial (cf. Coleman, 1990), and the inclusion of co-evolving network complicates matters further.

Theoretical tools that were particularly suited to deal with interdependencies of individual action and the transitions between different levels of aggregation, are game theory and agent-based modelling. In this paper, we follow an approach based on these tools to study effects of the macrolevel network structure on alcohol use. We aim to explain differences in average alcohol use between *groups* of adolescents (in this case, school classes) by means of the social network structure in the group at the start of the influence-selection process. We formulate the following research question:

To what extent can properties of the initial overall network structure explain differences in average alcohol use between school classes?

In answering this question, we use a theoretical model based on *strategic* interaction, in which we model social influence as a coordination game. In this model, we explicitly account for endogenous evolution of the network. Before we explain how our approach complements existing approaches to co-evolution of networks and behavior, we outline our theoretical model.

1.1 Coordination, influence, and alcohol use

In developing our model, we assume that when deciding whether to use alcohol, adolescents have incentives to choose the same behavior as the peers they interact with, for reasons that we outline below. This implies an interdependence between adolescents' decisions resembling the strategic structure of a *coordination game*. We assume that the outcome of this "game" determines the utility that a student derives from each of his friendship relations. The outcome of the game, in turn, depends on the behavioral choices of both students involved in a friendship relation. The game is shown in Figure 1, in general form and with numerical payoffs (only the relative values of the payoffs matter). In a friendship *network*, students play this game with multiple friends simultaneously. For ease of exposition, however, we first discuss the two-player setup and afterwards generalize this to a network setting.

Each player in the game has two options: either to drink or not to drink alcohol. This game is a coordination game because it has two Nash equilibria in which both players choose the same action. Thus, the players prefer to play the same action as their interaction partner, reflecting the basic idea that adolescents face a *pressure for conformity* in their behavior.

There are several reasons why such a peer influence might exist. First, there may be some intrinsic reason why an activity brings more utility if it is coordinated with others. A trivial example is games: playing ball is more fun if you can coordinate with someone to play with you. Similarly, drinking alcohol is most likely a social behavior, in the sense that the utility of use is higher if it is shared with someone else. A second incentive for coordination is *imitation*. During adolescence, people go through important changes, and consequently face many uncertainties. As a result, adolescents may look at their peers as a reference to help them determine which behavior is appropriate (Marsden and Friedkin, 1993). Third, there may be norms among groups of adolescents that promote conformity in general, also in the area of substance use (Sherif and Sherif, 1964). These three distinct pressures all lead to incentives for adolescents to coordinate their behavior. In the coordination game in Figure 1, this is represented by the fact that the payoffs for both players are higher when they choose the same action than when they choose different actions (i.e., $a > b$ and $d > c$).

The preference to coordinate does not imply that students are necessarily indifferent between using alcohol or not. In the structure of the game, we assume that students prefer abstinence to using alcohol, given that they coordinate their behavior. In Figure 1, this is reflected by the fact that $d > a$. This reflects the assumption that the disadvantages of using alcohol in terms of the financial costs, long-term health risks, and possible sanctions by parents and teachers are higher than the short-term gains.

Another feature of this game is that the "punishment" for failure of coordination differs between the actions. In our numerical example, if Player 1 chooses to drink and Player 2 chooses not to drink, the payoff of Player 1 is 8 while the payoff of Player 2 is 0. In other words, in a situation where one drinks and the other does not, this is worse for the one who does not drink. This assumption signifies that the use of alcohol has a number of effects on social behavior that have a negative impact on especially the social environment of the user, rather than on the user herself. For instance, drinking can lead to inflation of the ego, increased risk-taking, or downright aggression (Steele and Josephs, 1990). In situations where some use alcohol and others don't, such behaviors can have negative consequences especially for the non-users. In this sense, we can say that abstinence involves a higher risk on a lower payoff, such that the equilibrium in which both players drink can be classified in game- theoretic terms as a *risk-dominant* equilibrium (Harsanyi and Selten, 1988). In Figure 1, this feature is established by the assumption that $(a - b) > (d - c)$.

	Drink	Not drink
Drink	a, a	c, b
Not Drink	b, c	d, d

Drink	Not drink
$14, 14$	$8, 0$
$0, 8$	$20, 20$

Figure 1: Alcohol use as a coordination game, in general form and with numerical payoffs. $b < c < a < d$; $(a - b) > (d - c)$.

The game in Figure 1 provides a simple model for two actors in a friendship relation, deciding whether or not to drink. In reality, such choices take place in friendship *networks*, in which students maintain relations with several friends. We can extend our two-player model to a network model by assuming that every player plays with *several* interaction partners simultaneously. Each player can choose one action against all interaction partners (i.e., either consume alcohol or not), and receive the payoff as in Figure 1 from every interaction. Thus, students receive utility from every friendship relation separately, but must also adjust their behavior to several friends simultaneously.

A crucial assumption of this model is that actors cannot differentiate their behavior between different interaction partners: a student cannot choose to drink with one friend and not with the other. The rationale for this assumption is the idea that by choosing to drink in some situations, students make a general decision to "be a drinker," and thereby influence *all* their relations. Obviously, this is a simplification. In reality, it is very well conceivable that students behave differently with one friend than with the other. However, such differentiation is likely to be more costly in terms of effort than the situation in which one can simply use one mode of behavior. By considering alcohol use as an individual attribute rather than as relational choice, we also conform to the standard in the literature (e.g., Graham et al., 1991; Kirke, 2004; Bot et al., 2005), in which substance use is typically analyzed as an individual characteristic and not as a relational attribute.

In order to include selection, we assume that actors can choose with whom to interact. We assume that the coordination game is played repeatedly, and that in every iteration of the game, actors can choose both their behavior in the game, and with whom to play. When updating their behavior or their network links, actors choose the optimal response to what their interaction partners did in the previous interaction. Figure 2 illustrates this process. Actor i plays the coordination game with her two "neighbors" j and k. Actors i and j play one type of behavior, while k and l play the other type of behavior. At the same time, i and l have the opportunity to form a new link. We assume that maintaining social relations costs time and effort, such that every actor has to pay some cost for every link she maintains. The link between i and l is only formed if the expected benefits will outweigh the costs for both i and l.

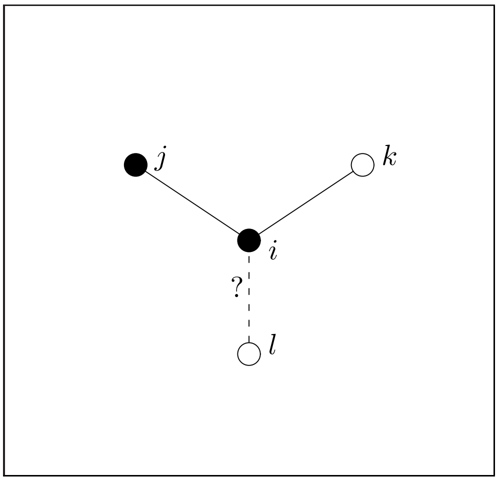

Figure 2: A coordination game in a dynamic network

Thus, we have sketched the outlines of a simple game-theoretic model for selection and influence in a dynamic network. We use a slightly more complex version of this basic model to derive specific hypotheses to be tested in the context of alcohol use. We first discuss how such a model could contribute to understanding influence and selection processes in comparison to earlier approaches.

1.2 Approaches to the study of selection and influence

Disentangling the simultaneous effects of influence and selection has been the focus of considerable research effort in the past decade. The basic problem was already recognized by Cohen (1977) and Kandel (1978), who noted that the effect of influence is likely to be overestimated if selection effects are ignored (the reverse also holds). A major breakthrough in the study of influence and selection was achieved by the introduction of actor-oriented, simulation-driven statistical models for longitudinal network data, implemented in the SIENA software (Snijders, 2001). In brief, this method works as follows: Given a set of subsequent observations of a social network, these panel data are considered to constitute "snapshots" of an underlying dynamic process. It is assumed that this process is driven by actors trying to optimize an *objective function* (roughly, a random utility function), both with regard to their own network position and with regard to their own behavior. The aim of the method is to estimate the components of this objective function based on the observed "snapshots" of this process. This is achieved by simulating the underlying process, and optimizing the fit between the simulated process and the observations based on maximum likelihood criteria. The effects of influence and selection can be identified as separate coefficients in the estimated objective function. Detailed expositions of the SIENA method can be found in Snijders (2006) and Snijders et al. (2007). Examples of applications of the method studying co-evolution of behavior and networks can be found in Steglich et al. (2006), Light and Dishion (2007), and Knecht (2008), among others.

Using the data we also analyze in this paper, Knecht (2008) applied SIENA to study the co-evolution of friendship networks and alcohol use to disentangle influence and selection effects. The results showed a clear selection effect (adolescents select friends based on similarity in drinking behavior), but provided only weak evidence for social influence.

In this paper, we take a somewhat different approach. We first use the coordination model to derive predictions about the relation between the initial state of a group (in this application, a school class) in terms of network structure and behavior and aggregate behavior at a later stage. We compute aggregate statistics on the groups in our data to create a dataset of *groups*, each observed at two different time-points. We then test statistically whether the groups in the data developed in the way that was predicted by the model.

To explain how this approach compares to SIENA, we highlight the most important differences and similarities. First and foremost, our method tests predictions at the *macrolevel* (that is, at the level of a whole network), while SIENA tests predictions at the individual level (namely, hypotheses about components of the utility function). In a sense, we can say that SIENA tests hypotheses on how the process works at the microlevel, while we test hypotheses on the *outcomes* at the macrolevel, assuming a specific theory on how the process works at the microlevel. We explicitly use initial states of the co-evolution process to predict outcomes, while SIENA merely conditions its simulations on the initial states. In this sense, the approaches are complementary. In the treatment of the data, the main difference is that we use only aggregate measures of network structure and behavior to test our predictions, while SIENA needs individual-level information.

A second difference between our approach and the SIENA approach is that rather than estimating properties of the network dynamics from the data, we assume a very *explicit* model of network formation and co-evolution. That is, we specify in detail the precise strategic nature of the interaction by means of the coordination game. In this respect, our model is more detailed in the specification of actors' incentives than SIENA. Similarly, we do not estimate the rate at which network changes can occur (the *rate function* in SIENA terminology), but instead make specific assumptions on this rate (in Buskens et al., 2008, we show that outcomes are robust under different assumptions on the speed of network dynamics).

Besides differences, there are also a number of similarities between our analysis and a typical SIENA analysis. First, and most importantly, both methods assume the possibility of an underlying co-evolution process in which both individual characteristics and the network change. Second, both methods rely on simulation to handle the complexity implied by a co-evolution process. Third, both methods assume that changes take place in "microsteps": only one link, or one actor's behavior, changes at a time.

By testing macrolevel hypotheses, we avoid two disadvantages of SIENA. The first problem concerns the theoretical interpretation of SIENA results; the second (related) problem concerns data requirements. Consider a co-evolution process based on coordination in a dynamic network. Suppose that, at some point, the network has reached a stable state in terms of behavior: no actor can improve her utility by changing her behavior, but some *can* improve their utility by changing ties. We observe the network for a few more "snapshots," in which very little change in behavior is observed (because it is already in, or close to, a stable state), but some tie changes are observed. SIENA results on these observations would probably indicate that there is no influence going on, but only selection. The theoretical implication of this would be that only selection plays a role in the considerations of the actors, even though stability in behavior is implied by a coordination game. In other words, social influence can only be identified as change, although theoretically, social influence could also be expressed by *stability*. The second problem is related to the first. Estimating effects of selection and influence in a SIENA model requires that enough changes in both the network and behavior are observed in the data. If an observed network is close to stability on one of the dimensions, this may lead to estimation problems in SIENA, forcing the analyst to drop these observations from the data (cf. Knecht, 2008, who can analyze only 78 out of 120 school classes).

The alternative approach we use here suffers less from these issues. Because our approach relies on examining *outcomes* of an underlying co-evolution process, it circumvents the problem of the interpretation of results that we identified above. Even if, in a given network, only relational changes are observed, our model still provides predictions that follow a model that assumes both selection and influence. Thus, also networks that are observed in a stable state can be compared with the predictions, because the theoretical model *predicts* what stable states should look like. As a result, also networks that are relatively stable contribute to the test of the hypotheses, while they would lead to estimation problems in SIENA. In the following, we apply our model to predict properties of stable states from initial conditions in terms of behavior and network structure. Results on this relation between initial conditions and outcomes can be used to identify selection- and influence effects. As we will show in the analysis, we find evidence that the emerging distribution of alcohol use is influenced by the *initial* network structure. We argue that this is a strong indication for the existence of influence: if the process would be driven only by selection, the network should only adapt to the distribution of behavior, and not vice versa. In principle, this logic could also be applied to identify selection effects: in that case, the model would predict an effect of the initial distribution of behavior on the emerging network structure. In this paper, however, we restrict the analysis to emerging behavior as the dependent variable because we are mainly interested in explaining differences in alcohol use.

Because our model predicts properties of stable states, our method does not require that enough changes in ties and behavior are observed in each network to make estimation possible. As a result, we are able to use a larger share of the available data to test our hypotheses. However, our approach *does* require that enough variation in initial conditions and outcomes is present in the data. In this study we meet this demand by using data on a large number of groups. Another disadvantage of our approach is that we cannot directly test hypotheses on individual decision processes, as is possible with SIENA.

Although our theoretical model takes into account that behavior and the network co-evolve, it provides the most informative predictions on the emerging distribution of *behavior*, and less in terms of the emerging *network*. The analyses in Buskens et al. (2008) show that the emerging network is almost perfectly determined by the emerging distribution of behavior: actors maintain only links with other actors with the same behavior. These predictions can be tested on the individual level, and therefore an empirical analysis of network formation based on our model would not provide addition insight as compared to a SIENA analysis. Moreover, from a substantive point of view, we are mainly interested in explaining differences in alcohol use, and less in explaining network structures.

2 Predictions

The theoretical model outlined above describes a dynamic process in which the network and behavior co-evolve. It can be shown analytically that this process may converge to a large variety of stable states (Jackson and Watts, 2002; Berninghaus and Vogt, 2006; Buskens et al., 2008). Thus, by itself, this model does not yet provide precise predictions on which stable states will occur. To obtain more informative predictions on which stable states are more likely to occur than others, one can use computer simulations. Buskens et al. (2008)

ran extensive computer simulations of the same model as we use here. The simulations resulted in a large dataset of initial conditions and resulting outcomes. This dataset was subsequently analyzed using conventional regression analysis methods, yielding predictions on how outcomes in terms of aggregate behavior depend on initial conditions, in terms of the initial distribution of behavior and the initial network structure.

We use these results to derive specific hypotheses on development of alcohol use among adolescents in school classes. To connect the simulation model with our empirical setting, we need to make a number of assumptions. We assume that alcohol use is for adolescents essentially a coordination game, as explained above: adolescents have incentives to display the same behavior as those they interact with; jointly using alcohol has (*ceteris paribus*) a lower utility than abstinence; and the risk involved in unilaterally using alcohol is lower than the risk of not unilateral abstinence. The network is the friendship network between adolescents in a school class. The underlying assumption is that the group of classmates constitutes a salient interaction context for adolescents. Because they spend a considerable share of their time at school among peers, we expect that they adjust their behavior to interaction partners from this group. This implies that we also assume that adolescents are not influenced by relations they might have *outside* their class. Admittedly, this is probably an unrealistic simplification of the situation. However, there is evidence that the students have most of their friendships and also their most important friendships in school classes (Knecht and Friemel, 2008).

We observe school classes at four different points in time. Applying our model, we aim to predict the behavior in the last observed period from the first observed period. Accordingly, all hypotheses are formulated in terms of effects of properties of a group (school class) at t_1 on properties of the group at t_4. A key group-level property at t_1 is the initial distribution of the *propensities* to choose alcohol use or abstinence. This propensity determines the likelihood that a student will use alcohol in the first "round of the game." The distribution of propensities determines only how the process starts; in subsequent time points, actions are exclusively the result of interaction in the coordination game.

As to effects of the initial network structure, we focus on the effects of network *density* and network *centralization*. In the simulation analyses by Buskens et al. (2008), these measures proved to have the largest effects on emerging behavior. *Density* refers to the extent to proportion of ties resent in the network, given the number of members of the network. *Centralization* is the extent to which ties are concentrated with relatively few individuals, rather than distributed uniformly among the network members (Snijders, 1981).

The first hypothesis serves as a "baseline" hypothesis, and relates the initial propensity to use alcohol to the resulting behavior at the end of the process:

Hypothesis 1. *The higher the average propensity to use alcohol in a class at t_1, the higher* *the proportion using alcohol at t_4.*

The next two hypotheses concern effects of the *initial density* of the network. Buskens et al. (2008) report that a higher initial density leads to a higher proportion of actors choosing the risk-dominant action, when starting from a situation in which the initial propensity is 50%. The intuition is that because the risk-dominant action is a stronger "attractor" (Young, 1998), more interaction at the start of the process (i.e., a higher density) leads more easily to convergence to the risk-dominant equilibrium. However, when the initial propensity is skewed, a higher initial density favors the action towards which the process already tended from the start.

This implies a hypothesis on a *main effect* of initial network density, and a hypothesis on an *interaction effect* between initial density and the initial propensity to use alcohol.

Hypothesis 2. *The higher the density of the network in a class with an equal distribution of initial propensity to use alcohol at t_1, the higher the proportion of students using alcohol at t_4.*

Hypothesis 3. *The higher the density of the network in a class at t_1, the stronger the effect of the proportion of students using alcohol at t_1 on alcohol use at t_4.*

We expect similar effects of *centralization* of the initial network as for density, but in the opposite direction:

Hypothesis 4. *The higher the centralization of the network in a class class with an equal distribution of initial propensity to use alcohol at t_1, the lower the proportion of students using alcohol at t_4.*

Hypothesis 5. *The higher the centralization of the network in a class at t_1, the weaker the effect of the propensity to use alcohol at t_1 on alcohol use at t_4.*

These two hypotheses signify that centralization of the network helps to counter the forces of the initial distribution of behavior and risk dominance. If the network is initially more centralized, there are actors in the network having relatively many interactions. Those actors are more influential, and if those actors happen to choose the risk-*dominated* behavior (i.e., *not* drinking), they are more likely to pull the rest of the network in this direction. Thereby, it is easier to "escape" from the "pull" of the risk-dominant behavior if the network is more centralized initially.

3 Data

3.1 Data collection

The data for this study were collected in a longitudinal survey project on 14 Dutch secondary schools, conducted in 2003 and 2004 (Knecht, 2006). From each school, all first-year classes were selected (between 5 and 14 classes per school, with an average of 9), and in each of these classes, all students were surveyed at regular intervals using written questionnaires. The first measurement took place shortly after the students entered the secondary school from primary education. The students were then surveyed again after three months, for a third time after another three months, and for a fourth and last time after another three months, resulting in a total of four waves. In total, 120 classes participated in all four waves. The survey included questions on personal characteristics of students, on various types of behavior (including alcohol use) and opinions, and various network measures.

The questionnaires were administered at school, with the students from each class together in a classroom. A researcher or research assistant was present at each session. Because the survey sessions were held during normal school hours, it could happen that not all students of one class were present. Moreover, some students may have joined a class between the moments of observation. Because of this, the number of students per class in the data may slightly differ between the different waves.

3.2 Variables and measures

3.2.1 Individual level measures

Personal networks Social relations in classes were measured using various *name generators*. In each wave, students were asked to name their best friends in class, the classmates with whom they spent leisure time, and those with whom they discussed personal matters. For each of these questions, they were allowed to name up to twelve classmates, using a list of codes for all classmates provided with the questionnaire. The maximum of twelve nominations was used only very rarely (up to 2% of the observations for any of the network measures), which indicates that this maximum was not a limitation on the measurement of nominations.

On the basis of these individual level measures, we construct networks at the class level. To verify that the results do not depend too much on the specific construction method, we use two different methods and report results using both methods.

For the first method, we combine measurements on three different name generators to identify interactions. We use nominations of “best friends,” spending leisure time together and discussing personal matters. We require that, *taken together*, nominations on the three variables are reciprocated. Thus, we assume that two students interact if they nominate each other, each on at least one of the three variables. This method takes into account that interpretations of friendly relations may differ between students: while student i might consider student j as one of her best friends, j might nominate i merely as someone with whom she discusses personal matters. As we are only interested in the extent to which students interact, we think that such mutual nominations, even though the interpretations of the relation slightly differ, can be interpreted as mutual interactions.

For the second method, we use only nominations from one name generator of best friends (as is most common in the network literature), and assume that if one student nominates another, the two interact. That is, we do not require that nominations are reciprocated, and we interpret every directed tie as a symmetric relation. Bilateral nominations are treated the same as unilateral nominations.

Both methods result, at the aggregate level, in a *non-directed* network in which all ties are bilateral. This is required to adequately test the predictions form our model, which explicitly assumes that adolescents have incentives to coordinate their behavior if they *interact*. By definition, interaction is non-directed, which implies a non-directed network.

Alcohol use The use of alcohol by the students was measured in different ways in the different waves of data collection. In the first wave, students were asked how often they used alcohol in the preceding three months. Answers could be given on a five-point scale: "never," "once," "2 to 4 times," "5 to 10 times," and "more than 10 times." In waves 2, 3, and 4, students were asked how often they had used alcohol in the preceding three months *with friends*, with the same answer categories as in the first wave. Thus, the measurement differs between the first wave and the other three waves in that the question of the first wave does not ask specifically about drinking with friends, but rather about drinking in general.

For this reason, Knecht (2008), who analyzes the same dataset, can use only the last three waves. Here we take the difference to be an advantage. In the context of our model, the measure in waves 2 to 4 represents alcohol use as far as it happens in a context of interaction with friends. This fits well in our theoretical framework, in which we assumed that alcohol use is a choice in a game of social interaction.[1] The measure used in wave 1, in contrast, we interpret as an indicator for the *individual propensity* to use alcohol before the influence/selection process that takes place among the students within one class starts. This corresponds with initial alcohol use at t_1 in our theoretical model. We feel that this is appropriate because the data collection in the first wave took place shortly after the start of the first year in secondary school. Thus, the three months mentioned in the question would refer for the largest part to the period just before the students entered secondary school, before they were influenced by the friendship network in their class. For this reason, it is not problematic that the measure of the first wave does not measure alcohol use with friends only.

Interpreting the measures in this way has the advantage that the complete time span between the four waves of data collection can be used in the analysis. The disadvantage is that the measure of wave 1 cannot be directly compared with the measures in the later waves. We can, however, use the measure of wave 1 as a predictor of alcohol use as measured in the later waves in a multivariate analysis (as we explain below).

3.2.2 Network level measures

Aggregate network measures Using the two operationalizations of interaction between students, we can construct a friendship network for each class at each time point. To be able to test our hypotheses, we compute network measures for each of these networks. The hypotheses are concerned with density and centralization. *Density* is defined as the number of existing ties divided by the number of possible ties, given the size of the network (Wasserman and Faust, 1994, p. 101). For *centralization*, we use the measure proposed by Snijders (1981), which is based on the (normalized) degree variance. Besides the measures needed for testing the hypotheses, we compute a measure of *relative network change* for descriptive purposes. This measure describes the extent to which the networks change between the different waves and is defined as the proportion of dyads in a network that have changed status (i.e., created a tie or deleted a tie) from one time point to the other. The measure is only defined as long as the set of nodes in the network does not change. Consequently, we cannot compute this measure for all networks on all time points. The number of networks for which we can compute the measure is at least 90 on each time point. Moreover, for the networks of "non-reciprocated friendship ties," (method 2) we also report the proportion of nominations in these networks that are actually reciprocated.

Aggregate measures of alcohol use We aggregate the individual measures of alcohol use into aggregate measures of alcohol use per class in two steps. First, we dichotomize the individual level measures between "1" (never) and "2" or higher (once or more). This is done to make the measures more consistent with the theoretical model, which assumes that only two different actions are possible. We choose this specific dichotomization because it is substantively clear: we now distinguish between those who do not drink at all and those who drink sometimes. Moreover, the empirical distribution on this variable is such that the large

[1] This assumption does not imply that we believe that there are no individual factors influencing alcohol use. In this study, however, we largely disregard these factors because we are interested in the *social dynamics* of alcohol use.

majority of the students does not drink. Taking this group as a distinct category therefore seems most appropriate. In the second step, we calculate the proportion of students drinking ("1" on the dichotomized variable) per class.

4 Methods of analysis

Our analytical strategy is set up as follows. We start the analysis with some descriptive statistics on the development of behavior and the network across the four waves. We then turn to regression analysis to test the hypotheses. In line with the analyses in Buskens et al. (2008) and Corten and Buskens (2010), we conduct an analysis at the macrolevel, using classes as the unit of analysis. The basic aim is to explain the level of alcohol use at the last observed time point, using measures characterizing the initial state per class as predictors. We use (linear) regression analysis, with the proportion of students that uses alcohol per class as the dependent variable. Because the classes were not independently sampled but are nested within schools, we use multilevel random intercept regression (Snijders and Bosker, 1999) with a random intercept at the school level.

Using a linear regression model to analyze a dependent variable that is a proportion is not without problems. We see a number of reasons why, in this case, using a linear model is not problematic. First, the distribution of our dependent variable does not show peaks at the edges of the distribution. In fact, the distribution closely resembles a normal distribution. Second, our models do not predict impossible outcomes, that is, values below 0 or higher than 1. Third, standard regression diagnostics do not indicate severe violations of model assumptions. Fourth, additional analyses (not reported here) using a logistic transformation of the original dependent variable do not lead to qualitatively different results. Buskens et al. (2008) and Corten and Buskens (2010) use logistic regression for grouped data for their analyses on the macrolevel. In those studies, use of logistic models was necessary because of the highly skewed distribution of the dependent variable. Because our dependent variable here is approximately normally distributed, we do not suffer from this problem. We instead prefer to use the somewhat simpler linear model, which allows for better treatment of the multilevel structure of the data. To verify whether the results are robust against different specifications of the network variables, we repeat the regression analyses for the two specifications discussed in Section 3.2.

To test hypotheses 3 and 5, we construct two interaction terms by multiplying the initial proportion drinking with density and centralization of the initial network, respectively. To facilitate the interpretation of respective main effects, we subtract 0.5 from the initial propensity before multiplication. This is necessary to test hypotheses 2 and 4, because these hypotheses predict effects of the network structure given that the initial propensity is 0.5. To ensure that the main effect of the initial propensity can be meaningfully interpreted, we center the values of initial density and centralization at their respective means before multiplication.

Thus, the interaction between the initial propensity to use alcohol and initial network density is computed as

$$\text{Interaction} = (\text{initial propensity alcohol use} - 0.5) \times (\text{density} - \text{mean(density)})$$

This construction ensures that, in the above case, the main effect of density in the regression equation can be meaningfully interpreted as referring to the situation in which the initial propensity is 0.5, while the main effect of the initial propensity can be interpreted as referring to the situation with average initial density.

5 Results

5.1 Descriptive results

Table 1 provides means and standard deviations on key measures at the individual level: the original five-point measure on drinking behavior, the dichotomized version of this measure, and the number of nominations of best friends, classmates with whom the respondent discusses personal issues, and classmates with whom the respondent spends leisure time. These three network variables are, at the individual level, *directed* measures; they measure the number of "outgoing" ties of a student.

The two measures of alcohol use show a rather consistent pattern: the average alcohol use decreases from wave 1 to wave 2, and then steadily increases. The initial decrease reflects the difference in measurement between wave 1 and 2 (see Section 3.2): because the initial measurement in wave 1 does not focus exclusively on alcohol use with friends but also captures drinking in other situations, the figures are somewhat higher. Overall, the averages are fairly low.

The average number of "best friends" nominations shows a slight increase over the first three waves, and then decreases again in the fourth wave. The other two network measures increase consistently over the four waves, but are clearly lower than the number of friends nominations.

In Table 2, we summarize the trends in measures on the aggregated (network-)level. Thus, in this table, the unit of analysis is the *class* rather than the individual student as in Table 1. The proportion of students drinking per class (as measured by the dichotomized variable) naturally shows the same pattern as the individual level statistics, but has a smaller standard deviation (the aggregate measure might be interpreted as a weighted average of the individual

Table 1: Descriptive statistics at the individual level, per wave

	Wave 1		Wave 2		Wave 3		Wave 4	
Variable	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Drinking (5 point scale)	1.854	(1.205)	1.459	(0.945)	1.636	(1.128)	1.756	(1.230)
Drinking (dichotomous)	0.424	(0.494)	0.242	(0.429)	0.301	(0.459)	0.342	(0.475)
"Friends" nominations	3.591	(2.585)	3.849	(2.743)	3.976	(2.793)	3.772	(2.638)
"Personal" nominations	1.225	(1.571)	1.676	(1.839)	1.882	(1.892)	1.899	(1.963)
"Leisure" nominations	1.466	(1.488)	2.006	(1.816)	2.320	(1.997)	2.444	(2.125)
Number of obs. (listwise)	2,826		2,723		2,768		2,781	

level measure). Both density and centralization are computed on the two different types of networks constructed by the methods described in Section 3.2.1. Averaged over classes, we do not see much of a trend in either of the two density measures. If anything, the figure shows a small increase over the waves 1 to 3, and a decrease in wave 4, which is consistent with the results at the individual level with regard to friends in Table 1. Centralization is rather low on both measures, and does not show much of a trend. Either network size is stable over time, and shows very little difference between the two methods. We report also the average change per class compared to the network in the preceding wave. The results on network change indicate that most of the changes occur between waves 1 and 2, and that the friendship network is slightly more dynamic than the network according to the combined measures.

In Table 3, we report the pairwise correlations between the dependent and independent variables to be used in the regression analyses. Because the two size measures are nearly identical, we report only results on the size of the combined networks. The results show that there are weak to modest significant correlations between the various network variables and across the behavioral variables. The fact that the measures for density and centralization are correlated between the different construction methods suggests that these construction methods do not lead to very different results. The correlation between density and centralization within each construction method is remarkably high, given that the measure for centralization is controlled for density. Closer inspection of the data shows that these correlations are caused by a small number of classes that have both high density and high centralization. Exclusion of these outliers, however, does not lead to different results of the regression analyses. Network size correlates significantly with all the other network measures, but not with behavior.

Table 2: Descriptive statistics at the class level, per wave (N=120)

	Wave 1		Wave 2		Wave 3		Wave 4	
Variable	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Prop. drinking	0.427	(0.133)	0.245	(0.120)	0.306	(0.136)	0.347	(0.151)
Combined network								
Density	0.092	(0.026)	0.103	(0.032)	0.110	(0.028)	0.107	(0.033)
Centralization	0.112	(0.054)	0.128	(0.056)	0.132	(0.057)	0.124	(0.053)
Network size	25.300	(4.076)	25.233	(4.140)	25.275	(4.122)	25.225	(4.103)
Change	-	-	0.094	(0.026)	0.082	(0.025)	0.079	(0.025)
Friendship network								
Density	0.202	(0.053)	0.226	(0.052)	0.232	(0.050)	0.222	(0.052)
Centralization	0.115	(0.073)	0.152	(0.071)	0.154	(0.064)	0.146	(0.059)
Network size	25.275	(4.079)	25.217	(4.141)	25.183	(4.095)	25.200	(4.095)
Change	-	-	0.167	(0.044)	0.133	(0.036)	0.127	(0.038)
Reciprocity	0.579	(0.084)	0.586	(0.090)	0.571	(0.090)	0.581	(0.086)

Table 3: Pairwise correlations between dependent and independent variables (N=120)

	1	2	3	4	5	6
1. Prop. drinking, Wave 1	-					
2. Prop. drinking, Wave 4	0.528**	-				
3. Density (combined)	0.100	0.098	-			
4. Centralization (combined)	0.089	0.233**	0.464**	-		
5. Density (friendship)	0.083	0.057	0.769**	0.522**	-	
6. Centralization (friendship)	0.028	0.103	0.493**	0.658**	0.633**	-
7. Network size	−0.034	−0.120	−0.591**	−0.593**	−0.628**	−0.544**

** $p < 0.05$

5.2 Multilevel regression using combined network measures

In this analysis, we use the networks as constructed by our first method, in which several types of name generators are combined. We conduct random intercept regression with average drinking behavior in wave 4 as the dependent variable, and class-level properties in wave 1 as predictors. We estimate three different models. In Model 1, we include only main effects of the initial proportion using alcohol, the initial density, and also control for network size. In Model 2, we add a term for the interaction between the two predictors. In Model 3, we add a main effect and interaction effect of centralization.

The results are displayed in Table 4. Model 1 shows a positive and strongly significant effect of drinking behavior in wave 1 as expected (Hypothesis 1), but no significant effect of initial density. We also find no significant effect of network size. In Model 2, the additional interaction effect is positive and significant, in accordance with Hypothesis 3. The main effect of initial density can in this model be interpreted as the effect of density for cases in which the initial proportion using alcohol is .5. This effect was expected to be positive (Hypothesis 2). The coefficient, however, is the opposite direction as expected but not significant. In Model 3, we add both the main effect of centralization and the interaction effect of centralization with drinking behavior in wave 1. Although both effects are in the expected direction (Hypotheses 4 and 5), they are not significant. Moreover, the likelihood ratio tests also indicates that although Model 2 is a significant improvement over Model 1, Model 3 does not further improve on Model 2. We therefore rely on Model 2, and conclude that only Hypotheses 1 and 3 are confirmed in this analysis.

5.3 Multilevel regression using non-reciprocated friendship ties

To examine to what extent our results depend on the specific network construction method, we repeat the analysis of the previous section using our second construction method, using (unreciprocated) friendship nominations as ties. Apart from this, the two analyses are identical.

Table 5 presents the results of this analysis. Overall, the results are consistent with the results in Table 4. We again find highly significant effects in the expected direction of the initial propensity and the interaction effect with initial density. Also, we again find no significant effects of initial density (main effect), centralization, or size, while the main effect of initial density is again in the opposite direction as expected.

We also find a number of differences as compared to the previous analysis. First, In Model 7, with effects of centralization included, the interaction term with density remains significant, in contrast with Model 3. A likelihood ratio test, however, indicates that Model

Table 4: Multilevel regression predicting average drinking behavior in wave 4 using combined network measures

		Model 1			Model 2			Model 3		
Variables	Prediction	Coeff.	SE	p	Coeff.	SE	p	Coeff.	SE	p
Prop. drinking, Wave 1	+	0.613***	(0.079)	0.000	0.609***	(0.077)	0.000	0.765**	(0.378)	0.043
Density	+	−0.833	(0.525)	0.112	−0.300	(0.557)	0.591	−0.296	(0.723)	0.683
Network size		−0.004	(0.004)	0.266	−0.003	(0.003)	0.333	−0.004	(0.004)	0.330
Drinking × Density	+				7.917***	(3.155)	0.012	6.575	(4.616)	0.154
Centralization	−							−0.038	(0.766)	0.961
Drinking × Centralization	+							2.134	(5.094)	0.675
Constant		0.257**	(0.127)	0.043	0.192	(0.126)	0.128	0.119	(0.221)	0.588
Variance components										
Var(Constant)		0.004	(0.002)		0.005	(0.002)		0.005	(0.002)	
Var(Residual)		0.012	(0.002)		0.011	(0.002)		0.011	(0.002)	
Log Likelihood		84.424962			87.442315			87.529721		
Likelihood ratio test: χ^2		48.20***			5.99**			0.22		
Number of observations		120			120			120		

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Likelihood ratio tests refer to the comparison with the previously estimated model; the likelihood ratio test for Model 5 refers to the comparison with the unrestricted model.

7 is no significant improvement over Model 6. Second, the effect of the interaction term with density in Model 7 is considerably smaller than in Model 2. Third, the main effect of centralization in Model 7 is positive, contrary to expectations, but not significant.

All in all, the results of this analysis are comparable to the results of the analysis with combined network measures. The likelihood ratio tests again indicate that Model 6, which includes the interaction term with density but no effects of centralization, is the preferred model. Most importantly, these results do not lead to different conclusions with regard to the hypotheses. The results suggest that the substantive conclusions are robust against the different specifications of the network variables.

5.4 Additional analyses

The analysis in Section 5.3 differs from the analysis in Section 5.2 in two respects: we changed from the combination of various network measures to friendship nominations only, and from reciprocated nominations to non-reciprocated nominations. For two combinations on these two dimensions, we found no substantive differences in results, but there are two combinations remaining: reciprocated friendship ties, and non-reciprocated combined ties. Results on these two combinations (not reported here) do also show no qualitative difference with the two analysis just discussed.

Our results consistently show an effect of the initial network structure, which leads to conclusions different from those by Knecht (2008) (see the concluding section). A possible explanation of this difference is the different use of data: because alcohol use was measured in a sightly different way in the first wave of data collection, Knecht (2008) could not use this part of the data and instead used the second wave as the first observation. In contrast, we interpreted the measure of alcohol use in the first wave as a pre-existing tendency for drinking, and used the first wave as the first time point. To check the validity of this explanation, we replicated the analyses of Section 5.2, using the data of the *second* wave of data collection. Although coefficients are in the same direction and of comparable magnitude, we find no significant results. Therefore, we cannot rule out that the differences in results are driven by a different choice of data, that is, by the difference in the measure of alcohol use or the fact that we look at a longer period.

6 Conclusions

In this paper we aimed to contribute to the understanding of selection and influence processes in the dynamics of alcohol use among adolescents. We did so by adopting a theoretical approach that interprets alcohol use as a coordination game in a dynamic network. Relying

Table 5: Multilevel regression predicting average drinking behavior in wave 4 using unreciprocated friendship ties

		Model 5			Model 6			Model 7		
Variables	Prediction	Coeff.	SE	p	Coeff.	SE	p	Coeff.	SE	p
Prop. drinking, Wave 1	+	0.610***	(0.080)	0.000	0.593***	(0.077)	0.000	0.591***	(0.076)	0.000
Density	+	−0.352	(0.257)	0.171	−0.094	(0.262)	0.719	−0.137	(0.288)	0.635
Network size		−0.004	(0.004)	0.311	−0.003	(0.003)	0.330	−0.002	(0.004)	0.494
Drinking × Density	+				4.518***	(1.540)	0.003	6.193***	(1.874)	0.001
Centralization	−							0.081	(0.203)	0.689
Drinking × Centralization	+							−1.972	(1.423)	0.166
Constant		0.246*	(0.132)	0.063	0.192	(0.129)	0.136	0.163	(0.129)	0.207
Variance components										
Var(Constant)		0.004	(0.017)		0.004	(0.002)		0.004	(0.002)	
Var(Residual)		0.012	(0.008)		0.011	(0.002)		0.011	(0.002)	
Log Likelihood		84.141062			88.204715			89.298062		
Likelihood ratio test: χ^2		47.63***			8.13**			2.9		
Number of observations		120			120			120		

* $p < 0.1$;** $p < 0.05$; *** $p < 0.01$. Likelihood ratio tests refer to the comparison with the previously estimated model; the likelihood ratio test for Model 1 refers to the comparison with the unrestricted model.

on simulation analyses of a game-theoretic model, we formulated hypotheses on the effects of initial conditions in terms of network structure and initial tendencies for alcohol use on resulting levels of alcohol use per school class. We tested these hypotheses using longitudinal data on alcohol use and social networks in Dutch high schools. Using various specifications of the independent variables, we were able to consistently confirm two hypotheses.

First, we find that the average initial propensity to use alcohol per class has a positive effect on average alcohol per class at a later stage. Second, we find that this effect becomes stronger, the higher the initial density of the social network in a school class. That is, in line with expectations, the density of the network *amplifies* the initial tendency of behavior.

We also predicted that initial network density should have a positive effect on alcohol use for classes that start with a average propensity to use alcohol of 0.5. However, we did not obtain any significant results on this hypothesis, and moreover, the estimated effect was consistently in the opposite direction than expected. Also, we were not able to obtain significant results on the effect of initial network *centralization*, where we predicted that centralization should have a negative main effect on alcohol use and should negatively interact with the initial behavioral tendency. While the direction of the estimated main effect is in some analyses opposite to the prediction, the interaction effect is in the expected direction.

How should our findings be interpreted? Although we were not able to test all hypotheses thoroughly, the fact that we found effects of the initial network structure on resulting behavior is revealing. The implication of this finding is that *influence* must play a role in the co-evolution process of alcohol use and network formation. If only selection would drive the process, then network formation would depend on the distribution of behavior, but not vice versa: the emerging behavior should be independent of the initial network. Instead, we find that there *is* an effect of the initial network. This conclusion contrasts with the findings of Knecht (2008), who found only evidence for selection using the same dataset. We return to this issue below.

The predictions on direction of the main effects of density and centralization depend on the assumption that using alcohol is the risk-dominant action in the coordination game. The facts that both effects were not significant, and that the main effect of density was consistently estimated in the opposite direction, suggest that this assumption might need to be revised. Such a revision might take two directions. First, the finding that the effect of density was consistently negative suggests that abstinence, rather than using alcohol, is the risk-dominant action. On the basis of this assumption, we would expect that a larger number of observations at the class level would yield a significant negative main effect of density. However, in that case we would also expect a significant *positive* effect of centralization, which we did not consistently find in our analyses. A second alternative assumption is that this coordination game is actually risk neutral, in the sense that neither of the two equilibria is risk dominant. In that case, we would indeed expect no main effects of density or centrality.

Such speculations, however, must be made with caution. The predictions on these main effects all refer to the situation in which the initial propensity is 50%. This means that the sizes and directions of these effects rely crucially on the exact definition of this majority, and therefore on the dichotomization of the dependent variable. While we think that our particular specification is well founded, arguments for different specifications are certainly conceivable. Therefore, one should be careful to draw strong conclusions based on the results on these main effects. Note, however, that the predictions on the *interaction effects* depend much less on the exact specification of the dependent variable.

Why do we find evidence for influence effects, while Knecht (2008) found only weak evidence? Part of the explanation could be that our approach indeed solves some of the problems of the SIENA approach applied by Knecht (2008), as we outlined in Section 1.2. That is, we are able to use also relatively stable classes for testing the hypotheses, and could therefore use more data. Besides the general methodological approach, however, there are some other differences between the two studies that could potentially account for the different findings. A first difference we already discussed concerns the use of data. Additional analyses showed that we cannot rule out this explanation of the differences in findings. Thus, it could be that the time from wave 2 to wave 4 is too short to observe network influence effects. Another possible source of that differences is that we were able to analyze a larger number of classes.

Second, the theoretical interpretation of influence is somewhat different between the two studies. Knecht (2008) analyzed the *directed* network, assuming that an adolescent is influenced by those peers she nominates as a friend. Thus, the influence is assumed to work in only one direction. In our model, we assume that influence takes place through *interaction*, which is by its nature undirected. While this difference is theoretically important, it should be noted that it is not exclusively a consequence of the different methodological approaches; a model with "two-way influence" would also be possible within the SIENA framework.

Third, partly as a consequence of the differing theoretical conceptualization of influence, our measures of the network are somewhat different. Knecht (2008) uses directed "best friend" nominations, while we use various combinations of network variables, including "best friend" nominations but also measures of spending leisure time together and discussing personal matters.

Given these differences, it is not clear how the differences in findings should be judged. We clearly encounter a discrepancy between findings at the macrolevel, where we do find evidence for influence effects, and the microlevel, where such effects could not be observed. This discrepancy poses a new puzzle that deserves more attention in future research on this topic. Such research should focus on the development of theoretical models that are consistent with empirical research on microlevel behavior and the macrolevel findings as presented in this study (see also Corten and Buskens, 2010). Overall, we conclude that the focus on effects of macrolevel network effects contributes to the explanation of emerging differences between classes and adds interesting new insights to the study of co-evolution processes.

References

Bauman, Karl E. and Susan T. Ennett. 1996. "On the Importance of Peer Influence for Adolescent Drug Use: Commonly Neglected Considerations." *Addiction* 91:185–198.

Bearman, Peter S., James Moody, and Katherine Stovel. 2004. "Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks." *American Journal of Sociology* 110:44–91.

Berninghaus, Siegfried and Bodo Vogt. 2006. "Network Formation in Symmetric 2×2 Games." *Homo Oeconomicus* 23:421–466.

Bot, Sander M., Rutger C. M. E. Engels, Ronald A. Knibbe, and Wim H. J. Meeus. 2005. "Friend's Drinking Behaviour and Adolescent Alcohol Consumption: The Moderating Role of Friendship Characteristics." *Addictive Behaviors* 30:929–947.

Buskens, Vincent, Rense Corten, and Jeroen Weesie. 2008. "Consent or Conflict: Coevolution of Coordination and Networks." *Journal of Peace Research* 45:205–222.

Centola, Damon and Michael W Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113:702–734.

Cohen, Jere M. 1977. "Sources of Peer Group Homogeneity." *Sociology of Education* 50:227–241.

Coleman, James S. 1961. *The Adolescent Society*. New York: The Free Press.

Coleman, James S. 1990. *Foundations of Social Theory*. Cambridge, MA: Belknap.

Corten, Rense and Vincent Buskens. 2010. "Co-evolution of conventions and networks: An experimental study." *Social Networks* 32:4–15.

Ennett, Susan T. and Karl E. Bauman. 1993. "Peer Group Structure and Adolescent Cigarette Smoking: A Social Network Analysis." *Journal of Health and Social Behavior* 34:226–236.

Epstein, Joyce Levy and Nancy Karweit (eds.). 1983. *Friends in School; Patters of Selection and Influence in Secondary Schools*. New York: Academic Press.

Giordano, Peggy C. 2003. "Relationships in Adolescence." *Annual Review of Sociology* 29:257–281.

Graham, John W., Gary Marks, and William B. Hansen. 1991. "Social Influence Processes Affecting Adolescent Substance Use." *Journal of Applied Psychology* 76:291–298.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78:1360–1380.

Harsanyi, John C. and Reinhard Selten. 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.

Hawkins, J. David, Richard F. Catalano, and Janet Y. Miller. 1996. "Risk and Protective Factors for Alcohol and Other Drug Problems in Adolescence and Early Adulthood: Implications for Substance Abuse Prevention." *Psychological Bulletin* 112:64–105.

Jackson, Matthew O. and Alison Watts. 2002. "On the Formation of Interaction Networks in Social Coordination Games." *Games and Economic Behavior* 41:265–291.

Kandel, Denise B. 1978. "Homophily, Selection, and Socialization in Adolescent Friendships." *American Journal of Sociology* 84:427–436.

Kirke, Deirdre M. 2004. "Chain Reactions in Adolescents Cigarette, Alcohol and Drug Use: Similarity through Peer Influence or the Patterning of Ties in Peer Networks?" *Social Networks* 26:3–28.

Knecht, Andrea B. 2006. *The Dynamics of Networks and Behavior in Early Adolescence [2003/04]*. Utrecht: Utrecht University.

Knecht, Andrea B. 2008. *Friendship Selection and Friends' Influence: Dynamics of Networks and Actor Attributes in Early Adolescence*. ICS dissertation series. Utrecht: Utrecht University.

Knecht, Andrea B. and Thomas N. Friemel. 2008. "Praktikable Vs. Tatsächliche Grenzen Von Sozialen Netzwerken: Eine Diskussion Zur Validität Von Schulklassen Als Komplette Netzwerke." Working paper, University of Erlangen-Nürnberg.

Lazarsfeld, Paul F. and Robert K. Merton. 1954. "Friendship as as Social Process: A Substantive and Methodological Analysis." In *Freedom and Control in Modern Society*, edited by Morroe Berger, Theodore Abel, and Charles H. Page, pp. 18–66. Princeton, NJ: Van Nostrand.

Light, John M. and Thomas J. Dishion. 2007. "Early Adolescent Antisocial Behavior and Peer Rejection: A Dynamic Test of a Developmental Process." *New Directions for Child and Adolescent Development* 118:77–89.

Marsden, Peter V. and Noah E. Friedkin. 1993. "Network Studies of Social Influence." *Sociological Methods and Research* 22:127–151.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27:415–444.

Moffitt, Terrie E. 1993. "Adolescence-Limited and Life-Course-Persistent Antisocial Behavior: A Developmental Taxonomy." *Psychological Review* 100:674–701.

Newcomb, Michael D. and Peter M. Bentler. 1989. "Substance Use and Abuse among Children and Teenagers." *American Psychologist* 44:242–248.

Sherif, Muzafer and Carolyn W. Sherif. 1964. *Reference Groups: Exploration into Conformity and Deviation of Adolescents*. New York: Harper and Row.

Snijders, Tom A. B. 1981. "The Degree Variance: An Index of Graph Heterogeneity." *Social Networks* 3:163–174.

Snijders, Tom A. B. 2001. "The Statistical Evaluation of Social Network Dynamics." In *Sociological Methodology*, edited by Michael E. Sobel and Mark P. Becker, volume 31, pp. 361–395. Boston, MA: Blackwell.

Snijders, Tom A. B. 2006. "Statistical Methods for Network Dynamics." In *Proceedings of the XLIII Scientific Meeting, Italian Statistical Society*, edited by S.R. Luchini, pp. 281–296. Padova: CLEUP.

Snijders, Tom A. B. and Roel J. Bosker. 1999. *Multilevel Analysis. An Introduction to Basic and Advanced Multilevel Modeling*. London: Sage.

Snijders, Tom A. B., Christian Steglich, and Michael Schweinberger. 2007. "Modelling the Co-evolution of Networks and Behavior." In *Longitudinal Models in the Behavioral and Related Sciences*, edited by Kees van Montfort, Han Oud, and Albert Sattora, pp. 41–71. Mahwah, NJ: Lawrence Erlbaum.

Steele, Claude M. and Robert A. Josephs. 1990. "Alcohol Myopia: Its Prized and Dangerous Effects." *American Psychologist* 45:921–933.

Steglich, Christian, Tom A. B. Snijders, and Patrick West. 2006. "Applying Siena: An Illustrative Analysis of the Co-evolution of Adolescents' Friendship Networks, Taste in Music, and Alcohol Consumption." *Methodology* 2:48–56.

Steinberg, Laurence and Amanda Sheffield Morris. 2001. "Adolescent Development." *Annual Review of Psychology* 52:83–110.

Swadi, Harith. 1999. "Individual Risk Factors for Adolescent Substance Use." *Drug and Alcohol Dependence* 55:209–224.

Wasserman, Stanley and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.

Watts, Duncan J. 2002. "A Simple Model of Global Cascades on Random Networks." *Proceedings of the National Academy of Sciences of the United States of America* 99:5766–5771.

Young, H. Peyton. 1998. *Individual Strategy and Social Structure. An Evolutionary Theory of Institutions*. Princeton, NJ: Princeton University Press.