The Effect of Density on the Level of Bias in the Network Autocorrelation Model[1]

Mark S. Mizruchi
University of Michigan

Eric J. Neuman
University of Michigan

ABSTRACT

Researchers interested in the effects of social network ties on behavior are increasingly turning to the network autocorrelation model, which allows for the simultaneous computation of individual-level and network-level effects. Earlier research, however, had pointed to the possibility that the maximum likelihood estimates used to compute the network autocorrelation model yielded negatively-biased parameter estimates. In this paper we use simulations to examine whether—and the conditions under which—a negative bias exists. We show that the network parameter estimate ρ is negatively biased under nearly all conditions, and that this bias becomes more severe at higher levels of both ρ and network density. We conclude by discussing the implications of these findings for researchers planning to use the network autocorrelation model.

---

[1] Research for the paper was supported by the College of Literature, Science, and the Arts and the Stephen M. Ross School of Business, both at the University of Michigan. We would like to thank Christopher Marquis for his significant contributions to this paper. We are also grateful to Phil Bonacich, Ron Breiger, Carter Butts, Malcolm Dow, Tony Smith, an anonymous reviewer, and Tom Snijders for their feedback on earlier drafts of the manuscript. Please address correspondence to Mark Mizruchi, Department of Sociology, University of Michigan, 500 South State St., Ann Arbor, MI, 48109-1382, phone: (734) 764-7444; FAX: (734) 763-6887; email: mizruchi@umich.edu.

As social network analysis continues to spread into an ever larger number of substantive areas, researchers have developed several ways to identify network effects. Because individual outcomes continue to be an important focus of social research, and because networks are by definition structures of relations rather than individuals, finding ways to determine the effects of these structures on individuals has been a vexing problem. Several solutions have been tried. Perhaps the most common is to use egocentric networks, which by definition are centered in the individual, and are thus amenable to the sampling of individual actors. Another common approach, used primarily when outcome variables are relational—such as similarity of behavior among actors—is to use dyads as units of analysis. In recent years, however, network analysts have developed an approach that allows for the computation of structural effects on individual outcomes. The technique that has been used, the network autocorrelation model, has its origins in geography, where it was used as a means of removing correlations among error terms across cases so as to render observations statistically independent (Cliff and Ord, 1981). As far back as the early 1980s, network analysts discovered that a variant of this model could be used as a means of detecting the influence of social ties. Instead of treating the autocorrelation across units as a nuisance to be removed, researchers could use it as a substantive effect to be measured and tested (Doreian, 1981; 1990).

Although the network autocorrelation model has gained an increasing number of adherents in recent years, a number of unresolved issues remain. Software remains relatively scarce and incapable of handling all but a few specific types of situations. Questions persist about the model itself, including the appropriate inputs into the network matrix (Leenders, 2002). There is also some indication that the model may contain a bias—that is, that the coefficients from the sampling distribution may systematically underestimate the true population parameter (Doreian, 1981; Dow, Burton, and White, 1982; Smith, 2004).

In this paper, we examine the question of whether the results of the network autocorrelation model are systematically biased. We also examine a potential source of this bias that has been overlooked in previous studies: the density of the network used as input to the model. Using a range of different networks of different sizes and with varying levels of network autocorrelation, we use simulations to examine the degree to which the density of the network affects the degree of bias in the parameter estimates derived from the model. We show that the model is biased at virtually all levels of autocorrelation, but that the bias becomes increasingly pronounced as network density increases. We show that this bias holds across a range of conditions, including both normalized and non-normalized networks. We conclude by discussing the implications of our findings for users of the model.

NETWORK AUTOCORRELATION MODELS

A fundamental assumption in regression analysis is the independence of observations. This is another way of saying that the error terms for the sampling distributions of each case are randomly distributed across cases. The violation of this assumption is common in time-series analysis, in which the same unit is observed in repeated cases. Whatever factors fail to explain a country's aggregate consumption in a particular year, for example, are likely to be correlated with the residuals from the previous or subsequent years. The independence assumption faces difficulties in certain kinds of cross-sectional analyses as well. Among the most common of these difficulties is the presence of correlated error terms due to geographic proximity. If one is examining rates of violent crime across U.S. states, for example, it seems likely that the unexplained outcomes in Mississippi and Alabama are more likely to be a result of similar forces than are those of Mississippi and California. The residual for Mississippi will therefore be more similar to that of Alabama than to that of California.

In the 1970s, geographers developed a series of techniques to address this problem of spatial dependence. In the standard OLS model, $Y = X\beta + \varepsilon$, where ε is a vector of randomly distributed disturbances. If ε cannot be assumed random, then some kind of transformation is necessary. Early work introduced two such transformations (Cliff and Ord, 1981; Ord, 1975). One way is to model the disturbances as being spatially autocorrelated by decomposing the vector ε into the form $\varepsilon = \rho W\varepsilon + \nu$, where W is an $N \times N$ matrix representing the distances between the observations, ρ is a parameter representing the degree to which the error terms are correlated, and ν is a vector of residuals that is now randomly distributed given the removal of the correlated errors into $\rho W\varepsilon$.[2] Anselin and Hudak (1992) refer to this as the spatial error model.

A second way of overcoming spatially dependent residuals in OLS is to model the autoregressive effects of the dependent variable: $Y = \rho WY + X\beta + \varepsilon$. As in the spatial error approach, W is an N x N matrix representing the distances between the observations while ρ represents the degree to which the observed values are correlated. Anselin and Hudak (1992) refer to this as the spatial lag model. Doreian (1980) terms it the spatial effects model. Cliff and Ord (1981) and Doreian (1980) stress that the two approaches are not competing models but rather reflect different assumptions about the underlying processes that drive the observed spatial autocorrelation.[3]

Subsequent to the development of these models, network theorists discovered that the spatial effects model could be extended to settings in which the autocorrelation stems from social, rather than purely physical, proximity (White, Burton, and Dow, 1981). In such cases, W becomes a matrix of social, as opposed to geographic, distances, and ρ is a substantive parameter estimating the extent to which an actor's outcome is affected by the behavior of those to whom he or she is socially proximate (Doreian, 1990). The spatial error and spatial effects models are often classified under the general heading of the network autocorrelation model. Our focus in this paper is on the spatial effects model.

This model has a major advantage over alternative network modeling approaches. To illustrate this, imagine that one is interested in the effects of proximity in a social network on an individual-level outcome variable. One way to model this association is through the use of dyads. Consider a case in which the dependent variable is the extent to which pairs of actors engage in similar behavior, and the investigator wants to know if the presence of a network connection between the members of a pair increases their likelihood of similar behavior. Assuming that we are interested in all possible pairs of relations within a network, this can be modeled by taking the $(N^2 - N) / 2$ dyadic relations

---

[2] The values in ν are sometimes referred to by econometricians as "innovations," to distinguish them from the residuals in the standard OLS model. See Greene (2000, pp. 749-750).

[3] See Leenders (2002, pp. 24-25) for a discussion of the different substantive interpretations that underlie the two models.

in the network and regressing the behavioral similarity of each dyad on the strength of the relation between them. As long as one controls for the non-independence of the observations due to the multiple appearance of the same actor in different dyads, this is a fully tractable problem. Techniques such as the use of actor-level fixed effects or quadratic assignment are readily available to handle the non-independence issue. The difficulty with this approach arises in cases where one must control for actor-level variables that might affect the size of the network effect. Converting individual-level variables to dyadic form is cumbersome, requiring the use of difference scores, with additional controls for baseline values. These difference scores are also difficult to interpret, which is especially a problem when they represent variables that might be of theoretical interest. Rather than rely on dyads, the network autocorrelation model allows the investigator to model an outcome variable for a single actor as a simultaneous consequence of both network and individual-level variables. This is a considerably clearer and more straightforward approach.

A dyadic approach may be equally, or even more, feasible if all of one's variables are relational (Mizruchi and Marquis, 2006). When the model includes a combination of network and individual-level variables, however, the network autocorrelation model appears to be superior. The model can handle any number of network parameters, in the form

$$Y = \rho_1 W_1 Y + \rho_2 W_2 Y + \ldots \rho_k W_k Y + X\beta + \varepsilon$$

where k equals the number of network parameters being estimated. See Dow (2007) for a detailed discussion of this model.

The dominant procedure for estimating ρ and β in the network (spatial) autocorrelation models is maximum likelihood (ML). Though other techniques such as Bayesian estimation (LeSage, 1997) and two-stage least squares (Anselin, 1990) have been proposed, an ML algorithm provided by Anselin (1988) serves as the basis for estimation routines in commonly used statistical packages such as R (Butts, 2006), Matlab (LeSage, 1999), and Stata (Pisati, 2001). The ML technique for estimating parameters in the network autocorrelation model was first introduced by Ord (1975), and explicated by Doreian (1981). We will not provide the details of the procedure here but instead refer the reader to those classic papers. We merely note that this technique assumes an underlying normal probability distribution function with normally distributed residuals.

NETWORK AUTOCORRELATION AS AN ESTIMATION APPROACH

In estimating a network autocorrelation model, we assume that the sample coefficient ρ represents an unbiased estimator of the true population parameter. We also assume that this property will hold regardless of the value of the parameter in the population. Dow, Burton, and White (1982), however, found, in data based on simulations, that the mean of ρ based on sampling distributions tended to be lower than the population ρ, and that the magnitude of the negative bias increased as ρ increased. Smith (2004) found a similar negative bias. These studies were concerned with the spatial error model rather than the network autocorrelation model and were thus estimating ρ as a nuisance term to be minimized rather than as a substantive predictor. Nevertheless, they do raise questions about whether the estimate of ρ in the network autocorrelation model might contain bias.

We came to this problem because of an interest in the association between network structure and the size of the network effect. Our concern was with whether there was a substantive association between network density and ρ. One argument, consistent with the work of Festinger, Schacter, and Back (1950), suggests that the extent to which actors are affected by that of their peers will be especially pronounced in highly cohesive groups. This is another way of saying that ρ will be positively associated with network density. This association, were it to exist, is a substantive one, however. As with any substantive finding, it depends on an unbiased estimate of the population parameter one is estimating.

Given the potential importance of network density on the effect of ρ, it is important to examine the extent to which ρ is unbiased across a range of density levels. In an earlier examination of randomly generated data (Mizruchi, Neuman, and Marquis, 2005), we found a tendency for the value of ρ to decline as density increased. Because the expected value of ρ with random data is 0, this suggested the possibility of a bias. In the following analysis we investigate whether ρ is biased, and if so, whether this bias varies across levels of network density.

TESTING PROCEDURE

To examine the degree of bias in the network autocorrelation model, we conducted a series of simulations, using combinations of ρ, the size of the network, the density of the network, and the number of variables in the X matrix. We began by selecting a particular value of population parameter ρ. We then constructed a series of networks of size N. One of the crucial issues in the use of network autocorrelation models is the content of W (Leenders, 2002). We chose the simplest possible version of W: a binary matrix, in which entries were coded as "1" if the two actors were directly tied, and "0" if they were not. To construct W, we set a target density of the matrix, ranging from .05 through .95, in increments of .05, such that the probability of any tie occurring was equal to the target density. We inserted a "1" when a tie was deemed to occur and a "0" otherwise, with "0" also appearing in the diagonal entries of W, as is standard in the use of this model.[4] We then normalized the rows so that they summed to 1, as recommended by Ord (1975), Cliff and Ord (1981), and Anselin (1988).[5]

We randomly set values of X plus the constant using a standard normal distribution. We then drew a random set of regression coefficients (β) for the X variables and the constant, also from a standard normal distribution. Finally, we randomly generated a vector of residuals, ε, using a normal distribution with a mean of 0. We varied the standard deviation, σ, across runs, however, in order to examine our results with different levels of "noise."

We generated "observed" values of Y by first rearranging the network autocorrelation equation from $Y = \rho WY + X\beta + \varepsilon$ to $Y = (I - \rho W)^{-1} (X\beta + \varepsilon)$ and then

---

[4] The results we report are all based on symmetric W matrices. We also ran simulations with an asymmetric W, but found no qualitative difference in the results.

[5] We examine the effects of using a raw (non-normalized) W matrix in a subsequent section of the paper.

inserting the assigned values of W, our X variables and their coefficients (β), our residuals (ε), and ρ. We then returned to the model and, using our observed values of Y, W, and X, we recomputed the equation, solving for the coefficients β and ρ.[6] If ρ is unbiased, then the value of ρ that we compute from these trials should have a sampling distribution with a mean approximately equal to the originally assigned value.

RESULTS

Our baseline model consisted of a network of size 50 (meaning that our W matrix was 50 x 50), three X variables and a constant (meaning that we estimate four β parameters), and a standard deviation of the residuals (σ) equal to 1 (meaning that ε was drawn from a standard normal distribution). For each of nine different values of ρ (-.8, -.5, -.2, -.1, 0, .1, .2, .5, and .8), we performed a trial, consisting of 100 replications of our baseline model, at each of 19 different target density levels, ranging from .05 through .95, in increments of .05. Figure 1 (panels a-i) presents scatterplots of the estimated ρ for each replication as a function of the actual density of W. Table 1 (panels a-i) presents the 95 percent confidence intervals around the sample mean of ρ for each trial.

A glance at the graphs in Figure 1 reveals a remarkable pattern. For each level of network autocorrelation, at lower densities, the mean sample values of ρ appear to be quite close to the population parameter. As the density moves past .5, however, the mean sample values of ρ begin to drop below the population value. By a density of .9, the mean sample values have taken a sharp negative turn and are now well below the population value.

FIGURE 1 ABOUT HERE

TABLE 1 ABOUT HERE

---

[6] Computation was performed in R using the "lnam" function in the sna package written by Carter T. Butts (2006). In other trials we have used routines written in Matlab (LeSage, 1999) and Stata (Pisati, 2001). Although we did not attempt to replicate every one of our analyses using these other software packages, the results of those that we did conduct were fully consistent with those reported here.

Table 1 displays the extent to which the sample means are significantly below what would be expected given the assigned population value of ρ. At each density level we computed a 95 percent confidence interval around the mean sample ρ. If the upper bound of the confidence interval falls below the assigned population value, we can conclude that the sampling distribution of ρ exhibits a significant negative bias.[7]

Consider the graph in panel a—the replications at ρ = -.8. From a density of .05 through .50, the assigned ρ of -.8 is within the bounds of every confidence interval, except at a density of .10, at which the confidence interval for the sample mean of ρ is actually above -.8.[8] At a density of .55, the sample mean slips significantly below the population parameter. Although the confidence interval of the sample mean barely straddles -.8 at densities of .60 and .70, the sample mean is significantly below the population value at every other point above .50. At a density of .95, the upper bound of the confidence interval is ρ = -1.94, far below the population value of -.8.

The results become considerably more pronounced as we move to higher population values of ρ. At -.5, the sample mean becomes significantly lower than the population ρ at a density of .25, and moves increasingly below -.5 thereafter, dropping sharply as the density approaches .9.[9] At -.2, the sample mean becomes significantly less than the population value at a density of .25, and remains significantly negative at every density level from .35 on. At population ρ levels of .5 and .8, the sample mean of ρ is significantly lower than the population ρ at every level of density, even .05. What is even more striking, however, is the extent to which the negative bias becomes stronger as density increases. At ρ = .5, the upper bound of the sample confidence interval is below .3 at every point from a density of .45. Even at a ρ as high as .8, the upper bound of the confidence interval is below .4 at every point from a density of .55.

---

[7] Because we computed 19 confidence intervals for each of the nine graphs (a total of 171 separate trials), the probability that any given sampling distribution will be significantly different from the population parameter is far greater than .05. As we will show, however, random sampling error cannot account for the consistently negative patterns that we observe over a broad range of conditions.

[8] Of the 171 trials under the conditions in Figure 1 and Table 1, this was the only case in which a sample mean of ρ was significantly above the assigned population value.

[9] The upper bound of the 95 percent confidence interval at a density of .25, rounded to four decimal places, is -.5001.

Of the 171 trials in this analysis, the mean sample ρ is significantly below the assigned population value in 139 cases. If we exclude ρ = -.8, then the sample mean is significantly below the population value in 132 of 152 cases. It is evident from this analysis that at least with a standard deviation of the residuals of 1 and a model that includes three X variables, the maximum likelihood approach to the network autocorrelation model contains a significant negative bias. Regardless of its level, the population parameter ρ tends to be systematically underestimated. Moreover, this negative bias increases as the density of the network increases. At high levels of density this negative bias is especially pronounced.

SENSITIVITY ANALYSIS

The previous analyses suggested the existence of a pronounced bias in network autocorrelation estimates. It is possible that the bias that we observed was specific to the conditions of our simulation, however. In this section of the paper we consider modifications of these conditions to see if the biases hold under alternative conditions.

Reducing the Model Noise

The first modification we make involves the variance of the residuals. We had set the standard deviation of the values in our ε vector equal to 1. In the second set of trials we maintained our model with three X variables and a constant, but reduced the standard deviation of the residuals to 0.5. This reduction should make it more difficult to observe a bias, since this model is designed, by definition, to provide a closer fit to the actual data. In this set of trials, we considered only six different values of population ρ: -.8, -.5, -.2, 0, .2, and .5.[10] As in the previous analysis, we performed 100 replications at each of 19 target density levels, ranging from .05 through .95. The results, for ρ values of 0 and

---

[10] As the variance of the residuals declined and ρ increased, the maximum likelihood algorithm that we used took increasingly long to converge. As a result, the full set of trials under such conditions could take two or more days to complete. Because the patterns in our findings seemed clear by the time we concluded the trials at ρ = .5, we elected not to run the trials at ρ = .8.

.5 only (to conserve space), are presented in panel a of Figures 2 and 3 and panel a of Tables 2 and 3. Results for all values of ρ appear in Figure S1 and Table S1 in the on-line supplement, at [insert link here].

FIGURE 2 ABOUT HERE

TABLE 2 ABOUT HERE

FIGURE 3 ABOUT HERE

TABLE 3 ABOUT HERE

As is evident from both the graphs and the confidence intervals in the tables, the results when we reduce the variance of the residuals (and thus improve the fit of the models) are virtually identical to those at the higher level of residual variance. At $\rho = 0$, the mean sample ρ is significantly below 0 when the target density is .1, and then at every density level greater than or equal to .3. At $\rho = .5$, the mean sample ρ is below .5 at every density level from .2 on. Of the 114 different conditions at which we produced trials, 78 exhibited a negative bias in ρ, in that the upper bound of the 95 percent confidence interval for the 100 replications was less than the population parameter. If we set aside $\rho = -.8$, which, consistent with the earlier analysis, yielded no bias at 12 of the 19 density levels, we observe a negative bias in 71 of 95 cases. Had we considered a ρ of .8, assuming that the previous pattern continued to hold, the prevalence of negative bias would have been even higher. In not a single one of the 114 trials did we observe a positive bias.

Can we make the bias disappear by shrinking the variance of the residuals even further? To address this question, we reduced the standard deviation (σ) of our residuals to .1. With a residual this small, the maximum likelihood algorithms we used took even longer to converge to a solution. In some cases our individual replications ran for as much as an hour before converging. Because of this difficulty, we conducted simulations on a network of size 40, with only 30 replications per density level rather than the 100 that we used in the first two experiments, and only 13 target density levels rather than the 19 we had used previously.

As is evident in our results, the use of this very small residual did allow us to remove most of the bias in our sampling distributions of ρ. As in the previous analysis (and subsequent ones), results for all of the values of ρ are available at [insert link here]. At $\rho = 0$ (panel b of Figure 2 and Table 2), the mean sample ρ is within the 95 percent confidence interval at every density except .9 and .95. At $\rho = .5$ (panel b of Figure 3 and Table 3), the mean sample ρ is within the confidence interval at every density level below .85. Among the 78 trials that we ran under these conditions, the 95 percent confidence interval was completely below the population ρ in only 10 cases, which is within the range of random chance given the number of trials. On the other hand, even under these restrictive conditions there was evidence of a negative bias. First, there was not a single instance in the 78 trials in which the confidence interval indicated a positive bias in ρ. Second, the means of the sample ρs were below the population value in 58 of the 78 trials ($\chi^2 = 18.51$ w. 1 df, $p < .001$), including in 35 of the 39 trials at ρs of 0 or greater. The fact that we conducted only 30 replications at each density level (rather than the 100 we had run in the previous two trials) contributed to the wider confidence intervals. We nevertheless can see some evidence of a negative bias at $\sigma = 0.1$.

Varying the Model

All of our network autocorrelation models to this point have involved an intercept and three exogenous X variables (in addition to the matrix W). To what extent would our results differ were we to vary the number of X variables? To examine this, we computed a model with an intercept but no X variables. In other words, the only causal variable in this model is the network matrix W. For these trials we returned to the situation in which we conducted 100 replications at each of 19 target density levels with a network size of 50, and we reset $\sigma = 1$. As with the most recent set of trials, however, we examined only six values of ρ (-.8, -.5, -.2, 0, .2, and .5). The results for ρ values of 0 and .5 are presented in panel c of Figures 2 and 3 and Tables 2 and 3.

As with our first two analyses, the findings based on a model with no X variables indicated a negative bias in ρ. The graphs indicate estimated values of ρ roughly at the population value from low to medium levels of density, but then a sharp negative turn at high density levels. Statistically, these results show the strongest negative bias among any of our analyses thus far. At ρ = 0, the mean sample ρ is significantly below 0 at every density level except .05. At ρ = .5, the mean sample ρ is below .5 at all 19 density levels. The upper bound of the confidence interval is below 0—that is, there is no evidence of a positive network effect—at every density level from .55 on, despite the true population ρ of .5. Among the 114 trials, in only 13 cases was there not a negative bias, and 10 of these occurred at the two lowest values of ρ (-.8 and -.5). For the remaining four levels of ρ, 73 of 76 trials exhibited a statistically significant negative bias. Moreover, in only 4 of the 114 trials was the mean sample ρ greater than the population ρ, and all four of these cases occurred at the two lowest values of ρ. In other words, the mean sample value was below the population parameter in 110 of 114 cases.

The fact that we observed an even stronger negative bias in our model with no X variables than we had in the model with three X variables led us to consider the possibility that we could reduce the negative bias in ρ by increasing the number of Xs. Panel d of Figures 2 and 3 and Tables 2 and 3 display the results of an analysis based on an equation with 40 X variables, at ρ values of 0 and .5 respectively. Although few regression models in the social sciences contain 40 parameters of substantive interest, it is not uncommon to compute equations with this number of variables, once controls such as fixed effects are included. For this analysis we preserved our σ of 1. Because of the processing time required to estimate these models, we considered only 13 target density levels at each ρ and conducted only 40 replications at each density level.

As the figures and tables indicate, the inclusion of 40 X variables does remove much of the negative bias in ρ. At ρ = 0, the mean sample ρ is significantly below 0 only at a density of .85 and above. At ρ = .5, the mean sample ρ is below .5 at a density of .4, as well as from .7 on. Among the 78 trials under these conditions, in only 15 did the full range of the 95 percent confidence interval fall below the population parameter. Still, there remained signs of a negative bias. First, in not a single case was the full range of the confidence interval above the population ρ. Second, in 62 of the 78 trials the sample mean was below the population ρ ($\chi^2 = 27.13$ w. 1 df, $p < .001$). Third, the frequency of negatively-biased ρ values increased at higher density levels and at higher values of ρ. At a ρ of .5, the full range of the confidence interval is below the population value at 6 of the 13 target density levels. Clearly, as in the case of the very low standard deviation of the residuals (panel b of Figures 2 and 3 and Tables 2 and 3), the extent of the negative bias of ρ is greatly reduced in the model with 40 X variables. In both cases, however, it took an extreme intervention to generate this condition, and in both cases evidence of at least a small negative bias remained.

Does the Size of the Network Matter?

Given the number of trials and replications we have reported to this point, we believe that we have established compelling evidence that the parameter ρ is systematically underestimated in the network autocorrelation model, and that this negative bias becomes more severe as both ρ and (especially) the density of the network increase. Almost all of our trials have been run on networks of size 50, however, while the remaining others have been run on networks of size 40. Is it possible that the negative bias would disappear, or at least decline, as the size of the network increases? To examine this, we reproduce our baseline model—with a σ of 1, three X variables, and an intercept—with a network of size 100. Because our simulations under these conditions took an extensive amount of time to converge, we ran only 50 replications at each of 13 target density levels, and we examined only three autocorrelation levels, -.5, 0, and .5. The results for the ρ values of 0 and .5 appear in panel e of Figures 2 and 3 and panel e of Tables 2 and 3.

As is evident from these results, the negative bias in ρ remains virtually as strong in the network of 100 as it did in the network of 50. At ρ = 0, the mean sample ρ is significantly below 0 at every density level from .3 on, the same as in the network of size 50. At ρ = .5, the mean sample ρ is below .5 at every density level from .15 on. Despite only 50 replications, the full range of the 95 percent confidence interval was below the population ρ in 28 of the 39 trials, including 20 of the 26 trials at ρ values of 0 and .5. In no case was the range of the interval above the population parameter. Moreover, the mean sample ρ was below the population parameter in 38 of the 39 trials. Only at the lowest density, .05, at the lowest ρ, -.5, was the mean sample value above the population parameter. In that case, the sample mean was -.499, or virtually identical to the population value of -.5. Meanwhile, as density increases, the sample values drop well below the population parameter.

Rather than declining in a larger network, then, the negative bias in ρ remains when we expand our network to a size of 100. Consider a comparison of the results from the network of 100 with the equivalent results from the network of 50: If we take only the trials with ρ values of -.5, 0, and .5, for the N of 50, a statistically significant negative bias (in which the full range of the confidence interval is below the population ρ) occurred in 50 of 57 trials, compared with 28 of 39 trials at N = 100. The sample mean was below the population parameter in 55 of 57 trials at N = 50, compared with 38 of 39 when N = 100. The difference in the prevalence of a significant negative bias (87.7 percent vs. 71.8 percent) is barely statistically significant ($\chi 2 = 3.86$ w. 1 df, $p < .05$).[11] It is not surprising, however, that the confidence intervals for the network of size 100 are wider than those of the network of size 50, given that we conducted only 50 replications at each target density level for the network of 100. The fact that in both cases virtually every one of the sample means was below the population parameter suggests that the negative bias in ρ is equally strong regardless of the size of the network.

Does Row Normalization of W Matter?

Authors of some of the classic discussions of spatial and network autocorrelation models have recommended normalization of the W matrix, such that the sum of each row equals 1 (Ord, 1975; Cliff and Ord, 1981; Anselin, 1988). This is typically done to facilitate the interpretability of the coefficients resulting from the model. Anselin (1988) has noted, however, that there is no mathematical or statistical requirement for doing so.[12]

---

[11] The .05 critical value of χ^2 with one degree of freedom is 3.84.

[12] As an example of how taken for granted it has been to normalize the matrix, Leenders (2002), in the most comprehensive discussion of the W matrix in the literature, discusses the issue at two points in the article, but only briefly mentions its benefits—in both cases in the paper's conclusion.

The results we have presented thus far have been based on a row-normalized W. This creates two conditions that could potentially affect our results. First, a normalized W is in nearly all cases asymmetric.[13] Second, although all network autocorrelation models will have values of ρ bounded by the reciprocal of the largest and smallest eigenvalues (Ord, 1975), the largest eigenvalue of a row-normalized W is always 1. This means that normalization creates a particular set of bounds for ρ. Because there is a strong positive association between the smallest eigenvalue and network density in a normalized matrix (Mizruchi, Neuman, and Marquis, 2005), this suggests the possibility that the relation between ρ and density reported above may have been affected by the normalization of W.[14]

To address this question, we ran simulations using our baseline parameters ($N = 50$, $\sigma = 1$, and a model with an intercept and three X variables) and a raw (non-normalized) W. With the raw W matrix, we can no longer use the same target values for ρ that we used in our previous simulations. As noted above, Ord (1975) has shown that the values of ρ are bounded by the reciprocal of the largest and smallest eigenvalues of W: $1/\lambda_{min} \leq \rho \leq 1/\lambda_{max}$. For any row-normalized W, $\lambda_{max} = 1$ and $\lambda_{min} > -1$. With a raw W, however, these specific constraints no longer hold. Smith (2007) has shown that a W with a density of 1 (which he terms W^*) has $\lambda_{max} = N - 1$, where N is the number of nodes in the network. In a fully-connected network of size 50, the maximum value of ρ is therefore 1/49, or .0204. When W is not fully-connected, its eigenvalues are determined by its particular configuration. Yet we know that for every such W the maximum eigenvalue will be less than $N - 1$, because $\lambda_{max}(W) < \lambda_{max}(W^*)$ for non-negative matrices $W < W^*$ (Horn & Johnson, 1985, Corollary 8.1.19). This implies that the upper bound of ρ for W matrices of size 50 is greater than the upper bound of ρ for W^* of size 50, which, as noted above, is .0204. If we limit ρ to be less than .0204, we satisfy the constraints of ρ for all densities of all raw Ws of size 50.

---

[13] The only condition in which a normalized W will be symmetric is one in which each tied pair of nodes has the same number of total ties, and thus the same row sum. In any single-component network with more than two nodes this means that all nodes must have the same number of ties.

[14] We thank Phil Bonacich for first raising the issue of the possible significance of whether W is normalized.

Adhering to this constraint, we conducted simulations with the non-normalized network at ρ values of 0, .005, .01, .015, and .02. We present results for three values of ρ: 0, .01, and .02. Results for all values of ρ are available in the on-line supplement [link]. As the results in Figure 4 and Table 4 indicate, the negative bias that we observed using the normalized W remains intact when we use the raw W. At ρ = 0, the mean sample ρ is significantly below 0 at every density level from .25 on. At ρ = .01, the mean sample ρ is significantly below .01 at every density level from .10 on. And at ρ = .02, the mean sample ρ is below .02 at every density level from .10 on.[15] As in our simulations with the normalized W, the negative bias in ρ becomes increasingly severe at higher levels of density.

FIGURE 4 ABOUT HERE

TABLE 4 ABOUT HERE

Because the values of ρ are not equivalent between the normalized and non-normalized W, we cannot directly compare the degree of negative bias between the two approaches. A look at the results in Figure 4 and Table 4 suggest, however, that the bias is neither more nor less severe when using raw W. For the row-normalized W and ρ = 0.5 (panel h of Figure 1 and Table 1), we find that the mean sample value for ρ is significantly less than the population parameter at every density level. The upper bound of the confidence intervals for these sample means is actually negative in four cases. The best comparison among the raw W simulations is ρ = .01 (panel b of Table 4). As is the ρ of .5 in the row-normalized W, .01 in the raw W is almost exactly half of the upper constraint on ρ (at least at high densities).[16] Here the mean sample value for ρ is significantly less than the population parameter for all but the lowest density (.05), and

---

[15] The upper bound of the 95 percent confidence interval at a density of .1, rounded to four decimal places, is .0197.

[16] To clarify, not only does the upper bound on ρ vary with network size when using raw W, it also varies with the density of W. We simplify matters here by keeping the target ρ constant across densities of a particular trial. In order to create simulations that are analogous to the row-normalized simulations, one should vary the target ρ according the density level of W. We have run such simulations and find a pattern of negative bias in ρ virtually identical to that found in our baseline model.

again there are four density levels for which the upper bound of the confidence interval is negative. Additional analyses comparing raw and row-normalized W matrices are available from the authors.

CONCLUSION

Our findings indicate that the network autocorrelation model produces parameter estimates of network effects that are systematically biased in a negative direction. This negative bias becomes increasingly pronounced as the density of the network increases: the average sample ρ turns sharply downward at high levels of density. These findings are invariant to the size of the network, and, except at very high numbers, are invariant to the number of parameters in the model. The negative bias can also be mitigated by significantly reducing the variance of the distribution of sample ρs at particular levels of density, but only by doing so at a level that is empirically unrealistic. Finally, the negative bias appears unrelated to the use of a row-normalized W or a raw W.

The fact that the negative bias in ρ becomes increasingly severe at high levels of network density means that when comparing network effects across networks, it is essential that researchers account for the density. Given the fact that the estimates of ρ tend to be biased downward, any findings of significantly positive ρs in substantive analyses should be taken as especially strong, particularly those based on high-density networks. A positive ρ under such circumstances would lend powerful support to the claim of Festinger, Schachter, and Back (1950), that pressures toward uniformity of behavior are especially pronounced in highly cohesive groups.

We found that whether a negative bias in ρ occurs at low density levels depends primarily on the level of ρ. In our baseline simulations, at density levels of .2 or lower, only one of 16 trials involving a negative population ρ yielded a significant negative bias in ρ, while 12 of the 16 trials involving a positive ρ yielded a significant negative bias. Still, the most substantial biases in estimates of ρ occur at high density levels, where, as we saw, the vast majority of trials yielded sample estimates below the population parameter. We should be attentive to the negative bias in ρ at both low and high density levels, but the bias at low density occurs primarily when ρ is positive. Given that network scholars are typically interested in cases with positive network autocorrelation, this increases the salience of our findings.

One potential limitation of this study is that the networks in our simulations are drawn from Bernoulli distributions. Such networks, based on random graph theory (Erdös & Renyi, 1959), tend to bear little resemblance to most social networks (Barabási, 2002; Robins and Morris, 2007). We have no reason to believe that the specific configuration of the network should affect the degree of bias in the estimation of ρ, except insofar as it affects the network density. We have, however, conducted preliminary simulations using small-world networks, which contain a specific structure, based on a particular path length and clustering coefficient (Watts, 1999). The findings thus far—available on request—are fully consistent with those presented above, based on random networks. Examination of the network autocorrelation model for particular network structures nevertheless remains a useful area for further work.

We close with two questions that flow from our findings, and that we believe need to be addressed: First, why does the network autocorrelation model produce a negative bias in ρ? Second, why does this negative bias become increasingly pronounced at high levels of both network density and ρ? Answers to these questions will have to await further investigation. In the meantime, we recommend that researchers using the network autocorrelation model be aware of the possible role of network density when interpreting their results. Our goal in future work will be to identify the characteristics of the network autocorrelation model that have led us to issue this recommendation. Our goal beyond that will be to provide a solution to the problem.

REFERENCES

Anselin, L., 1988. Spatial Econometrics: Methods and Models. Dordrecht: Kluwer Academic.

Anselin, L., 1990. Some robust approaches to testing and estimation in spatial econometrics. Regional Science and Urban Economics, 20: 141-163.

Anselin, L., Hudak, S., 1992. Spatial econometrics in practice: a review of three software options." Regional Science and Urban Economics 22, 509-536.

Barabási, A., 2002. Linked. Cambridge, MA: Perseus.

Butts, C.T., 2006. The sna package: tools for social network analysis, v. 1.4. Department of Sociology, University of California, Irvine.

Cliff, A.D., Ord, J.K., 1981. Spatial Processes: Models and Applications. London: Pion.

Doreian, P., 1980. Linear models with spatially distributed data: spatial disturbances or spatial effects? Sociological Methods and Research 9, 29-60.

Doreian, P., 1981. Estimating linear models with spatially distributed data. Sociological Methodology 11, 359-388.

Doreian, P., 1990. Network autocorrelation models: problems and prospects, in Griffith, I.D.A. (ed.), Spatial Statistics: Past, Present, and Future. Ann Arbor: Institute of Mathematical Geography, Monograph 12, pp. 369-389.

Dow, M.M., 2007. Galton's problem as multiple network autocorrelation effects: Cultural trait transmission and ecological constraint. Cross Cultural Research 41, 336-363.

Dow, M.M., Burton, M.L., White, D.R., 1982. Network autocorrelation: A simulation study of a foundational problem in regression and survey research. Social Networks 4, 169-200.

Festinger, L., Schachter, S., Back, K.W., 1950. Social Pressures in Informal Groups. Stanford: Stanford University Press.

Greene, W. H., 2000. Econometric Analysis. Fourth Edition. Upper Saddle River, NJ: Prentice Hall

Horn, R.A., Johnson, C.R., 1985. Matrix Analysis. Cambridge: Cambridge University Press.

Leenders, R.Th.A.J., 2002. Modeling social influence through network autocorrelation: constructing the weight matrix. Social Networks 24, 21-47.

LeSage, J.P., 1997. Bayesian estimation of spatial autoregressive models. International Regional Science Review, 20: 113-129.

LeSage, J.P., 1999. Spatial Econometrics. Web book of regional science, http://www.rri.wvu.edu/regscweb.htm.

Mizruchi, M.S., Marquis, C., 2006. Egocentric, sociocentric, or dyadic? identifying the appropriate level of analysis in the study of organizational networks. Social Networks 28:187-208.

Mizruchi, M.S., Neuman, E.J., Marquis, C., 2005. Does network structure affect the size of the network effect? The role of density in the network autocorrelation model. Paper presented at the Annual Meeting of the American Sociological Association. Philadelphia. August.

Ord, K., 1975. Estimation methods for models of spatial interaction. Journal of the American Statistical Association 70, 120-126.

Pisati, M., 2001. Tools for spatial data analysis. Stata Technical Bulletin 60, 21-37.

Robins, G., Morris, M., 2007. Advances in exponential random graph (p^*) models. Social Networks 29, 169-172.

Smith, T.E., 2004. Aggregation bias in maximum likelihood estimation of spatial autoregressive processes, in Getis, A., Mur, J., Zoller, H. (eds.), Spatial Econometrics and Spatial Statistics. New York: Macmillan, pp. 53-88.

Smith, T.E., 2007. Estimation bias in spatial models with strongly connected weight matrices. Unpublished manuscript. Department of Electrical Engineering, University of Pennsylvania.

Watts, D.J., 1999. Networks, dynamics, and the small world phenomenon. American Journal of Sociology 105, 493-527.

White, D.R., Burton, M.L., Dow, M.M., 1981. Sexual division of labor in African agriculture: A network autocorrelation analysis. American Anthropologist 83, 824-849.

Table 1. Sample mean and confidence interval of network effect (ρ) over 100 replications for each of 19 different target values (.05 to .95) of the density for W at 9 different population values of ρ. The model includes 3 X variables and a constant drawn from a standard normal distribution, a random row-normalized W of size 50, and error terms drawn from a standard normal distribution.

a. ρ = -0.8

density	mean ρ	confidence interval
.05	-0.809	(-0.845, -0.773)
.10	-0.774	(-0.792, -0.755)
.15	-0.779	(-0.807, -0.752)
.20	-0.799	(-0.835, -0.764)
.25	-0.765	(-0.806, -0.723)
.30	-0.759	(-0.815, -0.703)
.35	-0.816	(-0.879, -0.752)
.40	-0.768	(-0.826, -0.710)
.45	-0.806	(-0.876, -0.737)
.50	-0.851	(-0.927, -0.776)
.55	-0.927	(-1.003, -0.850)
.60	-0.889	(-0.982, -0.795)
.65	-1.057	(-1.153, -0.961)
.70	-0.915	(-1.035, -0.795)
.75	-1.029	(-1.176, -0.883)
.80	-1.152	(-1.301, -1.003)
.85	-1.384	(-1.543, -1.225)
.90	-1.736	(-1.924, -1.548)
.95	-2.180	(-2.421, -1.940)

b. ρ = -0.5

density	mean ρ	confidence interval
.05	-0.499	(-0.515, -0.483)
.10	-0.488	(-0.517, -0.458)
.15	-0.512	(-0.548, -0.477)
.20	-0.529	(-0.570, -0.488)
.25	-0.576	(-0.627, -0.526)
.30	-0.568	(-0.637, -0.500)
.35	-0.583	(-0.640, -0.525)
.40	-0.576	(-0.640, -0.513)
.45	-0.595	(-0.680, -0.510)
.50	-0.627	(-0.711, -0.544)
.55	-0.725	(-0.825, -0.626)
.60	-0.673	(-0.780, -0.566)
.65	-0.711	(-0.813, -0.609)
.70	-0.813	(-0.949, -0.678)
.75	-0.895	(-1.021, -0.768)
.80	-1.009	(-1.145, -0.873)
.85	-1.223	(-1.431, -1.015)
.90	-1.459	(-1.669, -1.248)
.95	-2.095	(-2.368, -1.823)

c. ρ = -0.2

density	mean ρ	confidence interval
.05	-0.200	(-0.220, -0.180)
.10	-0.195	(-0.223, -0.166)
.15	-0.248	(-0.288, -0.208)
.20	-0.247	(-0.295, -0.200)
.25	-0.288	(-0.339, -0.238)
.30	-0.231	(-0.290, -0.171)
.35	-0.277	(-0.343, -0.211)
.40	-0.324	(-0.410, -0.237)
.45	-0.311	(-0.381, -0.241)
.50	-0.385	(-0.471, -0.299)
.55	-0.355	(-0.451, -0.259)
.60	-0.365	(-0.465, -0.266)
.65	-0.547	(-0.659, -0.435)
.70	-0.650	(-0.755, -0.545)
.75	-0.572	(-0.695, -0.449)
.80	-0.791	(-0.929, -0.652)
.85	-0.966	(-1.138, -0.794)
.90	-1.249	(-1.463, -1.035)
.95	-2.062	(-2.337, -1.787)

d. ρ = -0.1

density	mean ρ	confidence interval
.05	-0.111	(-0.131, -0.091)
.10	-0.108	(-0.138, -0.078)
.15	-0.131	(-0.169, -0.094)
.20	-0.125	(-0.168, -0.082)
.25	-0.160	(-0.217, -0.103)
.30	-0.181	(-0.245, -0.117)
.35	-0.205	(-0.274, -0.1360)
.40	-0.228	(-0.297, -0.159)
.45	-0.198	(-0.271, -0.126)
.50	-0.267	(-0.353, -0.181)
.55	-0.402	(-0.495, -0.308)

e. ρ = 0

density	mean ρ	confidence interval
.05	-0.003	(-0.023, 0.017)
.10	-0.020	(-0.047, 0.008)
.15	-0.027	(-0.063, 0.009)
.20	-0.059	(-0.100, -0.018)
.25	-0.120	(-0.171, -0.069)
.30	-0.093	(-0.152, -0.034)
.35	-0.126	(-0.184, -0.067)
.40	-0.137	(-0.213, -0.062)
.45	-0.157	(-0.233, -0.081)
.50	-0.218	(-0.304, -0.131)
.55	-0.263	(-0.362, -0.164)

f. ρ = +0.1

density	mean ρ	confidence interval
.05	0.094	(0.069, 0.119)
.10	0.081	(0.054, 0.107)
.15	0.061	(0.023, 0.099)
.20	0.012	(-0.029, 0.053)
.25	0.057	(0.011, 0.102)
.30	-0.064	(-0.122, -0.007)
.35	-0.057	(-0.128, 0.015)
.40	-0.064	(-0.125, -0.003)
.45	-0.038	(-0.116, 0.039)
.50	-0.113	(-0.197, -0.029)
.55	-0.161	(-0.244, -0.078)

.60	-0.438	(-0.538, -0.338)
.65	-0.428	(-0.530, -0.325)
.70	-0.481	(-0.602, -0.361)
.75	-0.486	(-0.632, -0.340)
.80	-0.784	(-0.926, -0.642)
.85	-0.741	(-0.879, -0.603)
.90	-1.174	(-1.363, -0.986)
.95	-1.625	(-1.869, -1.381)

.60	-0.263	(-0.354, -0.172)
.65	-0.282	(-0.359, -0.204)
.70	-0.476	(-0.608, -0.345)
.75	-0.603	(-0.731, -0.476)
.80	-0.683	(-0.826, -0.539)
.85	-0.983	(-1.166, -0.800)
.90	-1.361	(-1.540, -1.182)
.95	-1.636	(-1.880, -1.392)

.60	-0.155	(-0.255, -0.055)
.65	-0.384	(-0.498, -0.269)
.70	-0.372	(-0.502, -0.243)
.75	-0.464	(-0.596, -0.332)
.80	-0.669	(-0.807, -0.532)
.85	-0.791	(-0.949, -0.633)
.90	-0.994	(-1.171, -0.817)
.95	-1.778	(-2.021, -1.520)

g. ρ = +0.2

density	mean ρ	confidence interval
.05	0.208	(0.186, 0.230)
.10	0.178	(0.145, 0.212)
.15	0.159	(0.120, 0.198)
.20	0.134	(0.087, 0.181)
.25	0.089	(0.036, 0.142)
.30	0.123	(0.073, 0.172)
.35	0.047	(-0.010, 0.104)
.40	0.029	(-0.031, 0.088)
.45	0.042	(-0.023, 0.108)
.50	-0.155	(-0.240, -0.069)
.55	-0.094	(-0.185, -0.004)
.60	-0.169	(-0.272, -0.066)
.65	-0.244	(-0.335, -0.152)
.70	-0.250	(-0.366, -0.133)
.75	-0.312	(-0.418, -0.205)
.80	-0.676	(-0.810, -0.543)
.85	-0.744	(-0.881, -0.607)
.90	-1.078	(-1.267, -0.889)
.95	-1.764	(-2.046, -1.482)

h. ρ = +0.5

density	mean ρ	confidence interval
.05	0.475	(0.456, 0.495)
.10	0.465	(0.440, 0.491)
.15	0.419	(0.383, 0.454)
.20	0.423	(0.385, 0.460)
.25	0.366	(0.313, 0.418)
.30	0.281	(0.232, 0.330)
.35	0.305	(0.246, 0.364)
.40	0.279	(0.201, 0.356)
.45	0.227	(0.163, 0.291)
.50	0.213	(0.136, 0.291)
.55	0.152	(0.073, 0.231)
.60	0.043	(-0.055, 0.141)
.65	0.043	(-0.064, 0.149)
.70	-0.086	(-0.182, 0.009)
.75	-0.090	(-0.204, 0.025)
.80	-0.324	(-0.474, -0.173)
.85	-0.562	(-0.736, -0.388)
.90	-0.854	(-1.040, -0.668)
.95	-1.599	(-1.842, -1.355)

i. ρ = +0.8

density	mean ρ	confidence interval
.05	0.784	(0.774, 0.795)
.10	0.753	(0.733, 0.773)
.15	0.683	(0.654, 0.712)
.20	0.636	(0.602, 0.671)
.25	0.601	(0.559, 0.644)
.30	0.571	(0.520, 0.622)
.35	0.530	(0.471, 0.588)
.40	0.444	(0.378, 0.511)
.45	0.412	(0.347, 0.477)
.50	0.354	(0.262, 0.446)
.55	0.291	(0.204, 0.378)
.60	0.301	(0.208, 0.394)
.65	0.165	(0.080, 0.250)
.70	0.105	(0.020, 0.189)
.75	-0.047	(-0.153, 0.060)
.80	-0.211	(-0.320, -0.102)
.85	-0.538	(-0.676, -0.400)
.90	-0.638	(-0.809, -0.467)
.95	-1.448	(-1.704, -1.192)

Table 2. Sample mean and confidence interval of network effect (ρ) by density target for 5 different variations of our baseline model, all with population ρ = 0. Unless noted, all simulations included 100 replications for each density target, 3 X variables plus an intercept, a random row-normalized W of size 50, and error terms drawn from a standard normal distribution (i.e., σ = 1.0). Nineteen different density targets (.05 to .95) were used in models a and c; 13 density targets were used in models b, d, and e.

Target ρ = 0

	a: σ =0.5		b: σ =0.1		c: intercept only		d: 40 X's		e: N=100	
dens.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.
.05	-0.005	(-0.019, 0.010)	-0.003	(-0.012, 0.005)	-0.014	(-0.047, 0.018)	-0.016	(-0.042, 0.010)	-0.001	(-0.035, 0.033)
.10	-0.028	(-0.050, -0.007)	-0.007	(-0.019, 0.006)	-0.081	(-0.126, -0.037)	0.033	(-0.003, 0.069)	-0.023	(-0.061, 0.016)
.15	-0.030	(-0.060, 0.001)	-0.007	(-0.018, 0.003)	-0.083	(-0.135, -0.031)	-0.001	(-0.045, 0.043)	-0.025	(-0.084, 0.035)
.20	0.012	(-0.021, 0.046)	-0.008	(-0.021, 0.005)	-0.088	(-0.149, -0.028)	-0.013	(-0.079, 0.052)	-0.021	(-0.084, 0.043)
.25	-0.037	(-0.077, 0.003)			-0.173	(-0.244, -0.103)				
.30	-0.056	(-0.095, -0.018)	-0.002	(-0.026, 0.022)	-0.218	(-0.291, -0.145)	-0.025	(-0.095, 0.045)	-0.083	(-0.155, -0.011)
.35	-0.063	(-0.103, -0.024)			-0.235	(-0.321, -0.150)				
.40	-0.079	(-0.128, -0.031)	0.015	(-0.014, 0.045)	-0.288	(-0.374, -0.202)	-0.000	(-0.087, 0.087)	-0.162	(-0.259, -0.064)
.45	-0.085	(-0.137, -0.033)			-0.323	(-0.419, -0.227)				
.50	-0.165	(-0.228, -0.102)	0.001	(-0.033, 0.034)	-0.336	(-0.456, -0.216)	0.056	(-0.053, 0.164)	-0.129	(-0.230, -0.027)
.55	-0.181	(-0.256, -0.106)			-0.402	(-0.510, -0.294)				
.60	-0.141	(-0.218, -0.065)	-0.005	(-0.079, 0.069)	-0.556	(-0.672, -0.441)	0.040	(-0.104, 0.183)	-0.255	(-0.421, -0.090)
.65	-0.205	(-0.283, -0.128)			-0.588	(-0.719, -0.458)				
.70	-0.241	(-0.324, -0.159)	-0.058	(-0.122, 0.005)	-0.836	(-0.978, -0.694)	-0.014	(-0.149, 0.121)	-0.429	(-0.588, -0.269)
.75	-0.287	(-0.384, -0.190)			-0.880	(-1.052, -0.707)				
.80	-0.367	(-0.481, -0.254)	-0.044	(-0.111, 0.023)	-0.947	(-1.113, -0.781)	-0.142	(-0.344, 0.061)	-0.536	(-0.726, -0.346)
.85	-0.506	(-0.645, -0.366)	-0.016	(-0.119, 0.086)	-1.416	(-1.611, -1.220)	-0.271	(-0.494, -0.049)	-0.899	(-1.157, -0.641)
.90	-0.623	(-0.777, -0.469)	-0.113	(-0.212, -0.014)	-1.662	(-1.897, -1.428)	-0.377	(-0.674, -0.081)	-1.159	(-1.416, -0.903)
.95	-1.192	(-1.430, -0.954)	-0.181	(-0.353, -0.008)	-2.905	(-3.190, -2.621)	-1.422	(-1.949, -0.894)	-1.883	(-2.210, -1.556)

a. Error terms drawn from a normal distribution with σ=0.5
b. Error terms drawn from a normal distribution with σ=0.1; 30 replications for each density target; matrix size of 40
c. No X variables; only an intercept
d. 40 X variables plus an intercept; 40 replications for each density target
e. Matrix size of 100; 50 replications for each density target

Table 3 Sample mean and confidence interval of network effect (ρ) by density target for 5 different variations of our baseline model, all with population ρ = 0.5. Unless noted, all simulations included 100 replications for each density target, 3 X variables plus an intercept, a random row-normalized W of size 50, and error terms drawn from a standard normal distribution (i.e., σ = 1.0). Nineteen different density targets (.05 to .95) were used in models a and c; 13 density targets were used in models b, d, and e.

Target ρ = +0.5

	a: σ =0.5		b: σ =0.1		c: intercept only		d: 40 X's		e: N=100	
dens.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.
.05	0.501	(0.490, 0.511)	0.498	(0.493, 0.503)	0.473	(0.446, 0.499)	0.509	(0.488, 0.530)	0.487	(0.466, 0.509)
.10	0.481	(0.461, 0.501)	0.504	(0.497, 0.511)	0.423	(0.389, 0.458)	0.493	(0.461, 0.524)	0.459	(0.417, 0.500)
.15	0.478	(0.453, 0.503)	0.500	(0.481, 0.518)	0.350	(0.301, 0.400)	0.458	(0.415, 0.501)	0.381	(0.319, 0.443)
.20	0.442	(0.412, 0.473)	0.496	(0.483, 0.508)	0.311	(0.255, 0.366)	0.506	(0.456, 0.556)	0.415	(0.363, 0.468)
.25	0.451	(0.417, 0.486)			0.310	(0.254, 0.367)				
.30	0.428	(0.392, 0.463)	0.499	(0.484, 0.515)	0.150	(0.078, 0.222)	0.427	(0.338, 0.516)	0.327	(0.237, 0.416)
.35	0.392	(0.349, 0.434)			0.108	(0.032, 0.184)				
.40	0.378	(0.325, 0.432)	0.495	(0.463, 0.527)	0.119	(0.041, 0.196)	0.385	(0.275, 0.496)	0.282	(0.189, 0.376)
.45	0.356	(0.310, 0.402)			-0.068	(-0.153, 0.018)				
.50	0.336	(0.279, 0.392)	0.489	(0.458, 0.520)	-0.036	(-0.122, 0.050)	0.397	(0.294, 0.500)	0.212	(0.124, 0.299)
.55	0.263	(0.202, 0.325)			-0.157	(-0.281, -0.034)				
.60	0.235	(0.158, 0.311)	0.479	(0.425, 0.533)	-0.249	(-0.344, -0.153)	0.415	(0.282, 0.549)	0.127	(-0.010, 0.264)
.65	0.202	(0.131, 0.273)			-0.364	(-0.484, -0.243)				
.70	0.220	(0.147, 0.292)	0.422	(0.330, 0.514)	-0.343	(-0.476, -0.209)	0.216	(0.067, 0.366)	-0.064	(-0.208, 0.080)
.75	0.087	(0.002, 0.172)			-0.605	(-0.763, -0.447)				
.80	-0.009	(-0.122, 0.104)	0.457	(0.376, 0.537)	-0.694	(-0.838, -0.550)	0.261	(0.118, 0.404)	-0.284	(-0.421, -0.146)
.85	-0.090	(-0.217, 0.037)	0.380	(0.294, 0.466)	-1.006	(-1.176, -0.835)	0.198	(-0.035, 0.431)	-0.523	(-0.716, -0.331)
.90	-0.344	(-0.476, -0.212)	0.386	(0.301, 0.470)	-1.396	(-1.594, -1.198)	-0.045	(-0.336, 0.245)	-0.777	(-1.021, -0.534)
.95	-0.946	(-1.174, -0.719)	0.342	(0.227, 0.457)	-2.273	(-2.543, -2.003)	-0.251	(-0.618, 0.116)	-1.432	(-1.775, -1.090)

a. Error terms drawn from a normal distribution with σ=0.5
b. Error terms drawn from a normal distribution with σ=0.1; 30 replications for each density target; matrix size of 40
c. No X variables; only an intercept
d. 40 X variables plus an intercept; 40 replications for each density target
e. Matrix size of 100; 50 replications for each density target

Table 4. Sample mean and confidence interval of network effect (ρ) over 100 replications for each of 19 different target values (.05 to .95) of the density for W at 2 different population values of ρ, 0 and .02. The model includes 3 X variables and a constant drawn from a standard normal distribution, a random raw (i.e., non-normalized) W of size 50, and error terms drawn from a standard normal distribution.

	a. Target ρ=0		b. Target ρ=0.01		c. Target ρ=0.02	
dens.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.	$\bar{\hat{\rho}}$	conf. int.
.05	0.001	(-0.007, 0.009)	0.003	(-0.006, 0.011)	0.013	(0.004, 0.022)
.10	0.002	(-0.004, 0.009)	0.001	(-0.004, 0.007)	0.014	(0.007, 0.020)
.15	-0.005	(-0.011, 0.000)	-0.000	(-0.005, 0.005)	0.011	(0.006, 0.015)
.20	-0.003	(-0.008, 0.002)	0.001	(-0.004, 0.006)	0.011	(0.006, 0.016)
.25	-0.007	(-0.011, -0.003)	-0.001	(-0.005, 0.003)	0.013	(0.009, 0.017)
.30	-0.008	(-0.011, -0.005)	0.002	(-0.001, 0.006)	0.014	(0.011, 0.017)
.35	-0.007	(-0.010, -0.004)	0.004	(0.000, 0.007)	0.011	(0.008, 0.014)
.40	-0.005	(-0.008, -0.001)	0.001	(-0.003, 0.004)	0.014	(0.010, 0.018)
.45	-0.008	(-0.011, -0.004)	0.000	(-0.003, 0.003)	0.011	(0.008, 0.015)
.50	-0.012	(-0.015, -0.008)	0.003	(-0.001, 0.006)	0.012	(0.009, 0.015)
.55	-0.008	(-0.012, -0.005)	-0.003	(-0.006, 0.000)	0.013	(0.011, 0.015)
.60	-0.010	(-0.012, -0.007)	0.001	(-0.002, 0.004)	0.010	(0.008, 0.013)
.65	-0.010	(-0.014, -0.007)	-0.001	(-0.004, 0.001)	0.008	(0.006, 0.011)
.70	-0.013	(-0.017, -0.010)	-0.002	(-0.005, 0.001)	0.010	(0.008, 0.012)
.75	-0.011	(-0.015, -0.008)	-0.005	(-0.008, -0.003)	0.006	(0.004, 0.008)
.80	-0.015	(-0.018, -0.011)	-0.003	(-0.006, 0.000)	0.008	(0.006, 0.010)
.85	-0.019	(-0.023, -0.015)	-0.005	(-0.008, -0.001)	0.006	(0.003, 0.008)
.90	-0.024	(-0.029, -0.019)	-0.015	(-0.019, -0.011)	0.008	(0.006, 0.010)
.95	-0.035	(-0.039, -0.030)	-0.025	(-0.029, -0.021)	0.005	(0.002, 0.008)

Figure 1. Scatterplot of the estimated network effect (ρ) against the actual density of W for each of 100 replications at 19 different target values (.05 to .95) of the density of W at 9 different population values of ρ. The model includes 3 X variables and a constant drawn from a standard normal distribution, a random row-normalized W of size 50, and error terms drawn from a standard normal distribution.

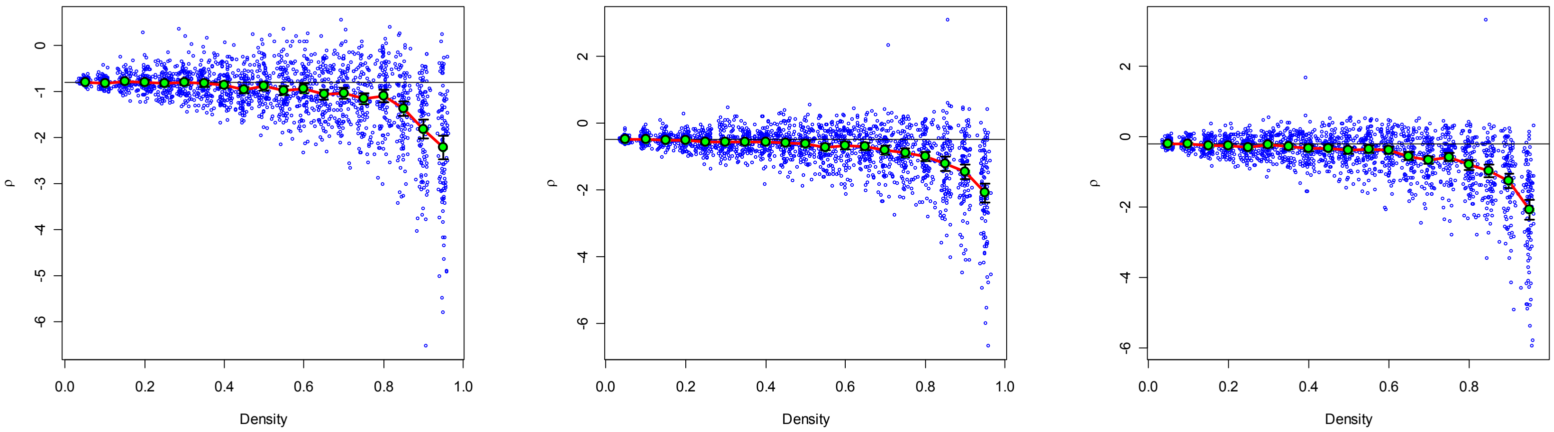

a. ρ = -0.8 **b. ρ = -0.5** **c. ρ = -0.2**

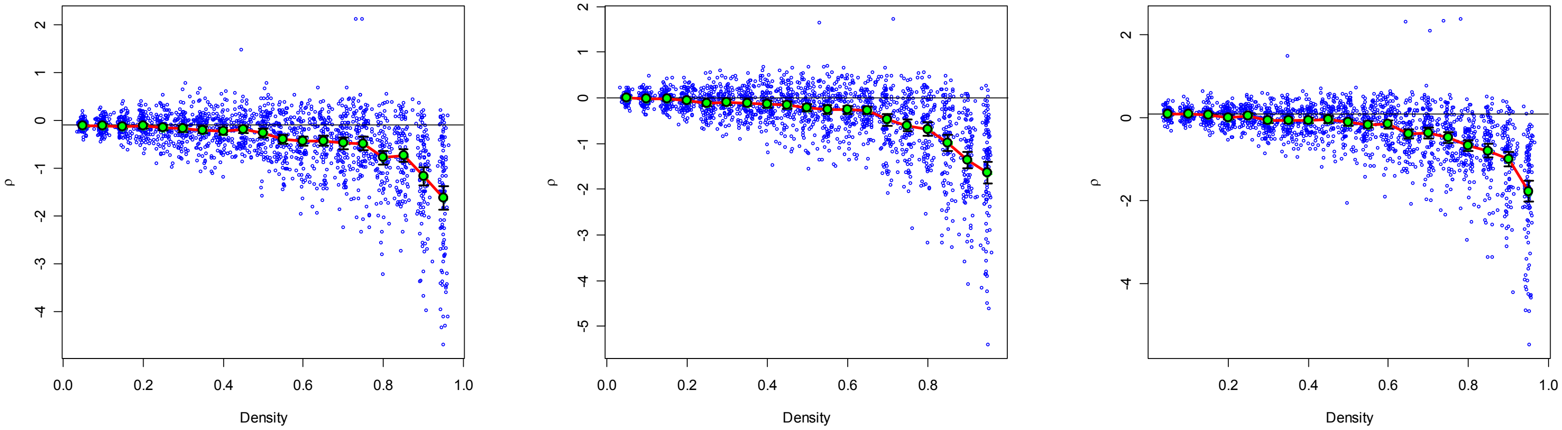

d. ρ = -0.1 **e. ρ = 0** **f. ρ = +0.1**

g. ρ = +0.2

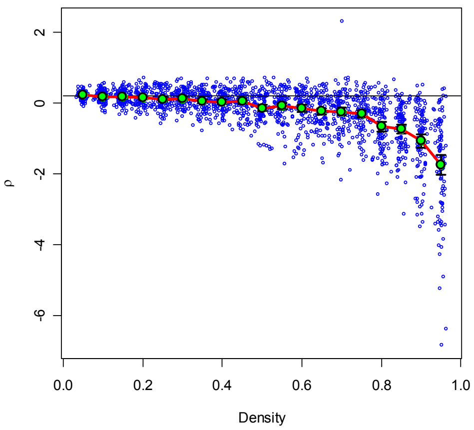

h. ρ = +0.5

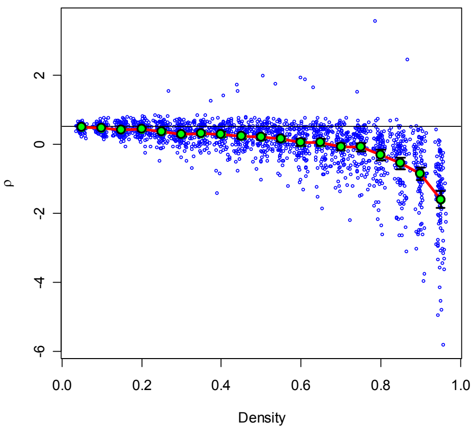

i. ρ = +0.8

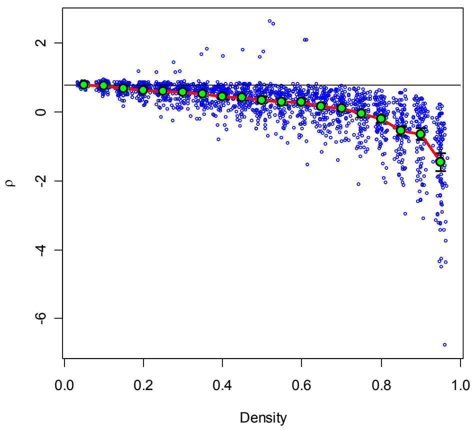

Figure 2. Scatterplot of the estimated network effect (ρ) against the actual density of W for 5 different variations of our baseline model, all with population ρ = 0. Unless noted, all simulations included 100 replications for each density target, 3 X variables plus an intercept, a random row-normalized W of size 50, and error terms drawn from a standard normal distribution (i.e., σ = 1.0). Nineteen different density targets (.05 to .95) were used in models a and c; 13 density targets were used in models b, d, and e.

a. σ = 0.5 **b. σ = 0.1** **c. intercept only**

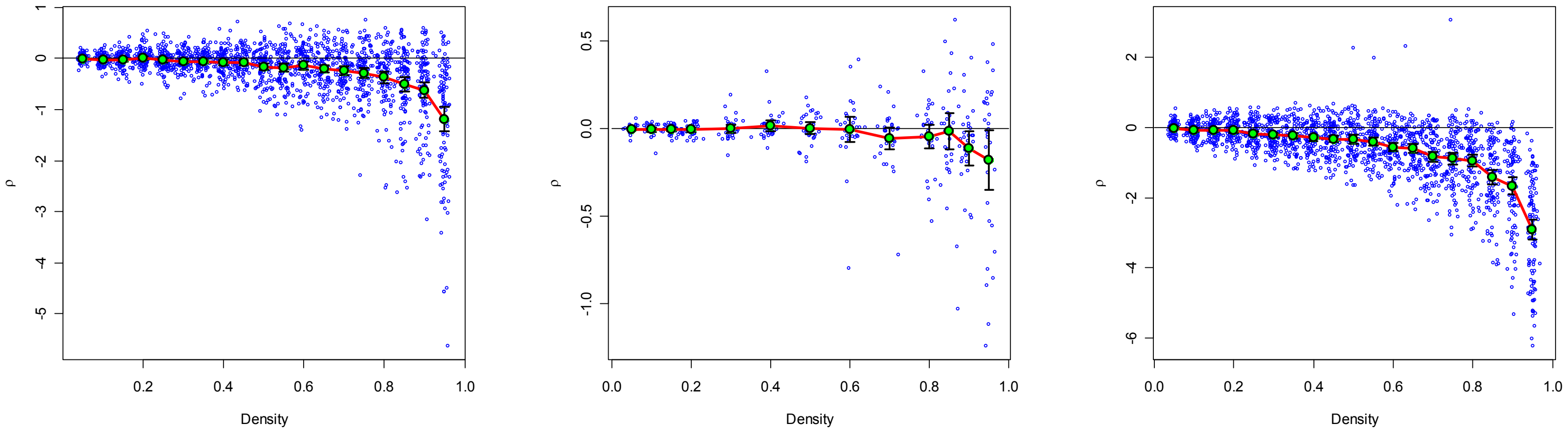

d. 40 X's **e. N = 100**

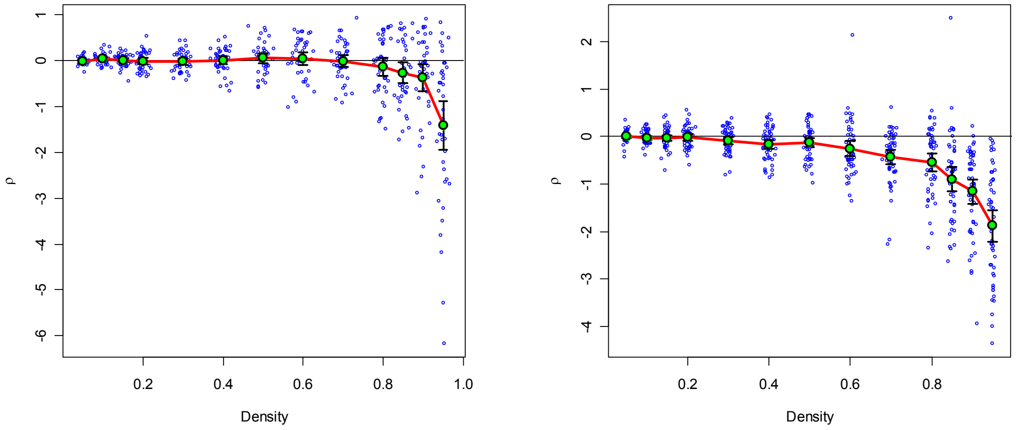

a. Error terms drawn from a normal distribution with σ=0.5
b. Error terms drawn from a normal distribution with σ=0.1; 30 replications for each density target; matrix size of 40
c. No X variables; only an intercept
d. 40 X variables plus an intercept; 40 replications for each density target
e. Matrix size of 100; 50 replications for each density target

Figure 3. Scatterplot of the estimated network effect (ρ) against the actual density of W for 5 different variations of our baseline model, all with population ρ = 0.5. Unless noted, all simulations included 100 replications for each density target, 3 X variables plus an intercept, a random row-normalized W of size 50, and error terms drawn from a standard normal distribution (i.e., σ = 1.0). Nineteen different density targets (.05 to .95) were used in models a and c; 13 density targets were used in models b, d, and e.

a. σ = 0.5 **b. σ = 0.1** **c. intercept only**

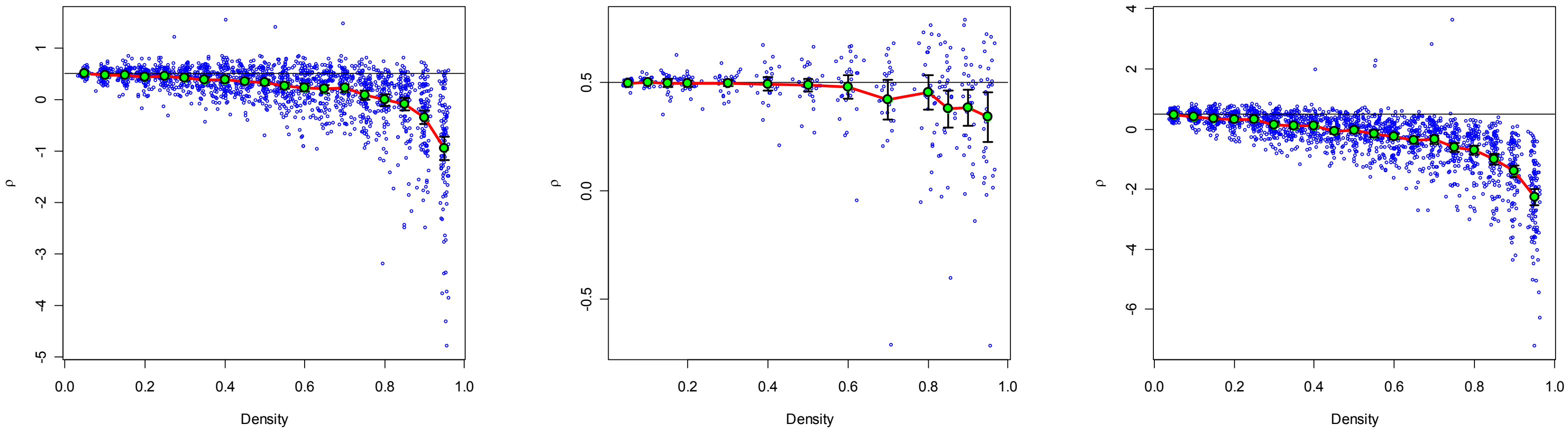

d. 40 X's **e. N = 100**

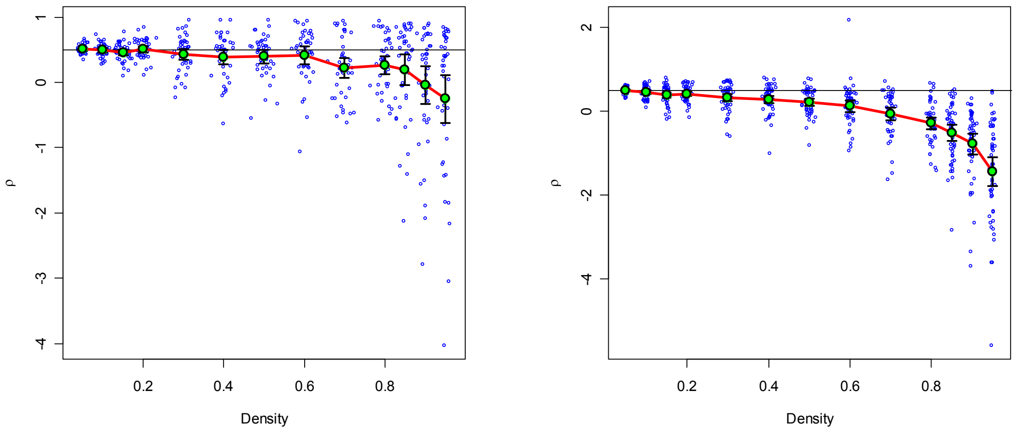

a. Error terms drawn from a normal distribution with σ=0.5
b. Error terms drawn from a normal distribution with σ=0.1; 30 replications for each density target; matrix size of 40
c. No X variables; only an intercept
d. 40 X variables plus an intercept; 40 replications for each density target
e. Matrix size of 100; 50 replications for each density target

Figure 4. Scatterplot of the estimated network effect (ρ) against the actual density of W for each of 100 replications for each of 19 different target values (.05 to .95) of the density for W at 2 different population values of ρ, 0 and .02. The model includes 3 X variables and a constant drawn from a standard normal distribution, a random raw (i.e., non-normalized) W of size 50, and error terms drawn from a standard normal distribution.

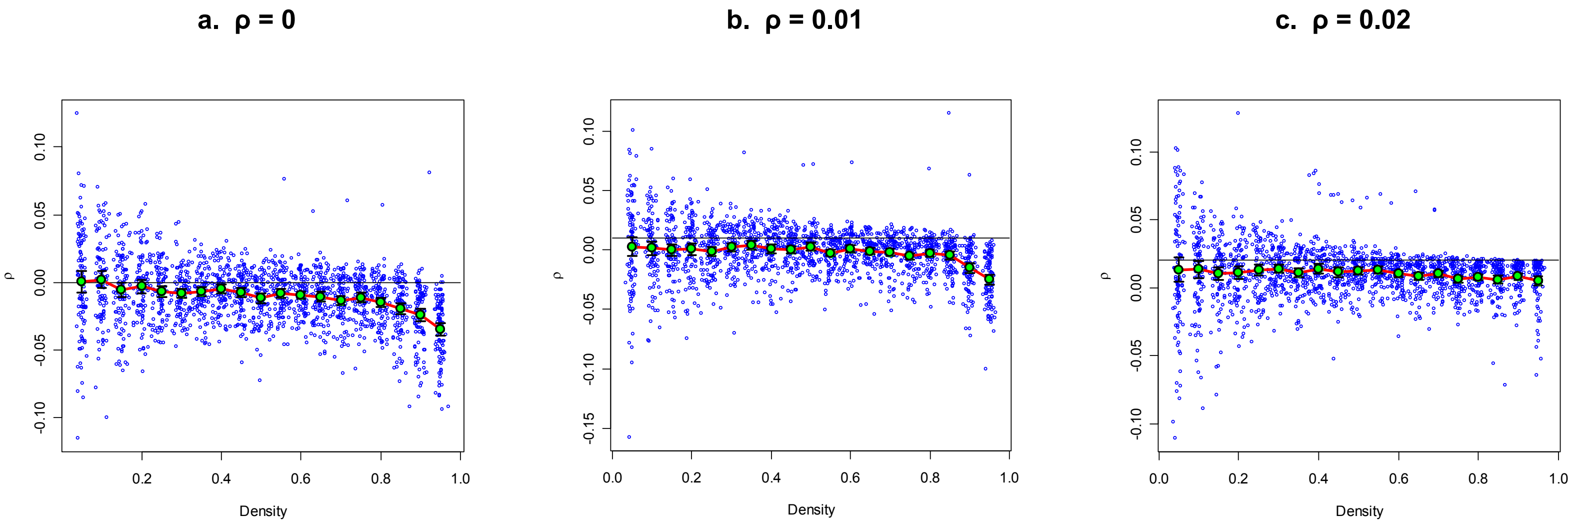

a. ρ = 0

b. ρ = 0.01

c. ρ = 0.02