For network effects models of network autocorrelation we use Monte Carlo methods to study the relative properties of three estimation methods. The methods are the iterative maximum likelihood estimation (Ord, 1975; Doreian, 1981), ordinary least squares, and a regression-based "quick and dirty" substitution for iterative MLE. Of the three, OLS is clearly the inferior estimation method and MLE the superior method. We recommend the use of the maximum likelihood method when network autocorrelation models are to be estimated.

Network Autocorrelation Models

Some Monte Carlo Results

PATRICK DOREIAN
KLAUS TEUTER
CHI-HSEIN WANG
University of Pittsburgh

Our concern focuses on some of the technical issues arising from the use of linear regression models when the data points are interdependent. A cursory examination of the mainstream journals in most social sciences reveals the widespread use of regression methods. Among the assumptions, usually implicit, underlying the use of these models is one concerning the independence of data points.[1] However, many situations exist in which the data points are not truly independent. Doreian (1981) provides a partial listing of examples in which the interdepen-

---

EDITOR'S NOTE: *Patrick Doreian served as editor to prepare a special issue of* Sociological Methods & Research *on autocorrelation in sociological research. Enough papers to fill an entire issue were not forthcoming. However, this article and the one following were prepared for the special issue, and the* SMR *Board of Editors recommended publication of these two. I thank Professor Doreian for the effort he made, and it should be recognized that his work commitment was nearly as great as it would have been had a full issue resulted.*

AUTHORS' NOTE: *We appreciate greatly the comments of an anonymous reviewer, which improved the manuscript considerably.*

SOCIOLOGICAL METHODS & RESEARCH, Vol. 13 No. 2, November 1984 155-200

dence stems from contiguity or adjacency in geographical space. White et al. (1981) deal with a network autocorrelation formulation of Galton's problem in which the dependency is generated by diffusion or a cultural splintering process. Loftin (1972) provides another example of Galton's problem viewed in terms of spatial autocorrelation.

Researchers who employ regression models in their research are typically concerned with two problems: (1) assessing whether or not a variable is significant, and (2) determining the values of coefficients viewed as parameters. Some evidence suggests that the pursuit of both objectives is seriously compromised by the use of ordinary least squares (OLS) methods without taking into account the interdependence among the data points. Loftin and Ward (1983) provide a telling example in which the presumed relation between population density and fertility is not supported when autocorrelation is considered although OLS applied to the same data set permits the inference that fertility does depend on population density.[2] White et al. (1981) provide examples where the specific parameter estimates differ when spatial autocorrelation is ignored compared to when it is explicitly considered.

Thus, instances in the literature alert us to the *possibility* that we may be misled by our OLS tools when there are interdependencies between the data points of a particular data set. However, these instances do not, by themselves, provide any indication of how serious the risk is either for making incorrect inferences or reporting inaccurate parameter estimates for a given linear model. Some analytical results and some Monte Carlo studies provide complementary pieces of evidence by which we obtain some indication of the properties of estimation procedures for linear models when we are confronted by interdependent data points.

THE NATURE OF THE MODEL

In the context of network[3] autocorrelation, broadly construed, two kinds of models are distinguished: the network effects model and the network disturbances model. A discussion of these two types of models is found in Doreian (1980). Our concern here is with the network effects model and we note that a simulation study has been conducted for the network disturbances model (Dow et al., 1982).

The linear model conventionally estimated is as follows:

$$y = X\beta + \epsilon\,;\ \epsilon \sim N(0, \omega I) \qquad [1]$$

When the network effects[4] model is considered, it is usually represented as

$$y = \rho Wy + X\beta + \epsilon : \epsilon \sim N(0, \omega I) \qquad [2]$$

Obviously, if $\rho = 0$, then the model in equation 2 reduces to that in equation 1. Thus, for the purposes of this article, equation 2 will be taken to be the model requiring estimation, with equation 1 being a special case. Given such a model, we will explore the properties of three estimation strategies: (1) ordinary least squares applied to the special case; (2) the maximum likelihood method (as described in more detail in the following section); and (3) a "quick and dirty" attempt to estimate equation 2 directly by ordinary least squares. This method is also discussed in the following section.

ESTIMATION STRATEGIES

If the model is given in equation 2 then, following Ord (1975), a maximum likelihood approach can be undertaken. Assuming equation 2, the likelihood function to be maximized is (with $A = I - \rho W$)

$$\ell(y) = \text{const} - (N/2)\ln\omega + \ln|A| + (1/2\omega)(y'A'Ay - 2\beta'X'Ay + \beta'X'X\beta) \qquad [3]$$

We estimate ρ as the scalar that minimizes

$$(-2/N) \sum_{i=1}^{N} \ln (1 - \rho\lambda_i) + \ln (y'My - 2\rho y'MWy + \rho^2 (Wy)'MWy) \qquad [4]$$

where $M = I - X(X'X)^{-1}X'$ and the λ_i are the eigenvalues of W.

Once ρ is obtained, $\hat{\beta}$ is given by:

$$\hat{\beta} = (X'X)^{-1} X'Ay \text{ and } \hat{\omega} \text{ by}^5 \qquad [5]$$

$$\hat{\omega} = \left(\frac{1}{N - K - 1}\right) (y'A'MAy) \qquad [6]$$

with K being the number of regressors. The residual $\hat{\epsilon}$ is obtained as $\hat{\epsilon} = (I - \hat{\rho}W)y - X\hat{\beta}$. If OLS is used directly on equation 1, the assumption is $\rho = 0$, with equation 3 reducing to the form of the log-likelihood function for the classical normal linear regression model. Then equation 5 reduces to

$$\hat{\beta} = (X'X)^{-1}X'y \qquad [7]$$

and equation 6 reduces to

$$\hat{\omega} = (1/N-K-1)y'My = (1/N-K-1)\ (\hat{\epsilon}'\hat{\epsilon}) \qquad [8]$$

Where $\hat{\epsilon}$ is the residual from the regression.

Conventional ordinary least squares is particularly easy, whereas the maximum likelihood approach is more computationally burdensome. The maximum of the log-likelihood function given in equation 3 can be found from a direct search procedure (see Doreian, 1981) or it can be done numerically via an interative algorithm (see Ord, 1975). The latter is used here.

Doreian and Hummon (1976: 138-140), in a discussion of spatial models, suggest the direct estimate of equation 2 by ordinary least squares with Wy as another variable entered on the right side of the equation and view such an approach as a variant of generalized least squares. This we term a "quick and dirty" estimation method. Although the parameter estimates returned by this method will be inconsistent, it is reasonable to ask whether the strategy is useful either for providing some guidance concerning the autocorrelation structure of the data or for providing parameter estimates that are sufficiently good for practical purposes.

DESIGN OF THE MONTE CARLO STUDY

Data were generated by specific regimes of parameters and analyzed by all three estimation strategies. The parameters are the slope and intercept coefficients, the variance of the disturbance term, ω, and the value of the network autocorrelation parameter, ρ. A given matrix, W, representing the interdependencies between specific data points and a matrix of observations on the independent variables, X, were taken and remain fixed throughout all of the simulation. In fact, W is taken to be the spatial adjacency matrix used by Doreian (1981) and the independent variables matrix is constituted by a column of 1s for the constant, and two of the independent variables for the Louisiana data set[6] reported in Table 2 of Doreian (1981). The vector of coefficients $(\beta_0, \beta_1, \beta_2) = (10, 0.3, -0.3)$ is used in the first set of simulations. This will be referred to as the "full model," although at a later point we report on simulations where, first, β_1 is set to 0 and, second, β_2 is set to 0. Given our interest in equation 2, batches of 100 random normal disturbance terms, ϵ, were generated with mean 0 and variance ω and then were fed through the filter implied by the following equation:

$$y = (I-\rho W)^{-1} X\beta + (I-W)^{-1}\, \epsilon \qquad [9]$$

With the specification of a regime, the fixed W and X, and the generation of the random disturbance terms, everything on the right-hand side of equation 9 is given. In each of 100 runs under a given regime, y and X together with W constitute the data set analyzed by ordinary least squares (OLS), maximum likelihood (MLE), and quick and dirty (QAD) methods. For OLS, ρ is assumed as 0 while MLE and QAD attempt to estimate it. Under all methods there are attempts to estimate β and ω.

At issue were the following:

(1) the bias, if any, of parameter estimates
(2) the true sample variability of estimates
(3) the reported sample variability (standard error)
(4) the frequency of incorrect and correct inferences

The spatial parameter will be bounded by -1 and $+1$. However, as virtually all empirical instances show positive values of ρ, if they show any value at all, the Monte Carlo studies have ρ confined to $0 \leq \rho \leq 0.9$ as the computation involved in the simulation is very large. Further, because of the computational burden, we consider only 0.1, 0.3, 0.5, 0.7, and 0.9 as values of ρ and only 4 values for the variance of the disturbance term, namely, 25, 49, 81, and 121[7] were used.

The data reported in the simulation results presented below include the following:

(1) The mean value of each coefficient estimate (as a simulation approximation for the expected value of that parameter estimate).
(2) The actual standard deviation of these estimates (as a simulation approximation to their true variation). This we label SDE.
(3) The mean of the reported standard error for each estimate (as a simulation approximation to the theoretical value of the standard error). This we label MRSE (not to be confused with the root mean squared error—RMSE).
(4) The range of the parameter estimates.
(5) A count of both the correct and incorrect inference decisions.

These inference decisions are of two kinds: for any parameter, denoted by θ, a null hypothesis of (1) H_0: $\theta = 0$, and (2) H_0: $\theta = \theta_{gen}$, where θ_{gen} is the actual parameter value used in the data generation.

ON THE BIAS AND SAMPLE VARIABILITY OF OLS

The use of simulation is unnecessary in addressing the first issue listed in the previous section when OLS is concerned. The OLS $\hat{\beta}$ is biased. From equation 7,

$$\begin{aligned} \beta &= (X'X)^{-1}X'y \\ &= (X'X)^{-1}X'(A^{-1}X\beta + A^{-1}\epsilon) \\ &= (X'X)^{-1}X'A^{-1}X\beta + (X'X)^{-1}A^{-1}\epsilon \end{aligned} \qquad [10]$$

So

$$E\hat{\beta} = (X'X)^{-1}X'A^{-1}X\beta \qquad [11]$$

If $\rho = 0$, then $E\beta = \hat{\beta}$ as required. But for $\rho \neq 0$, $(X'X)^{-1}X'A^{-1}X$ is not the identity matrix and β is biased for all regimes of spatial autocorrelation. The magnitude of the bias is also obtained easily. From equation 11 we have

$$\begin{aligned} E\hat{\beta} &= (X'X)^{-1}X'(I - \rho W)^{-1}X \\ &= (X'X)^{-1}X'\left(\sum_{i=0}^{\infty} \rho^i W^i\right)X\beta \\ &= (X'X)^{-1}(X'X)\beta + (X'X)^{-1}X'\left(\sum_{i=1}^{\infty} \rho W^i\right)X\beta \\ &= \beta + (X'X)^{-1}X'\left(\sum_{i=1}^{\infty} \rho^i W^i\right)X\beta \end{aligned} \qquad [12]$$

Hence, the bias is given by

$$\text{Bias} = (X'X)^{-1}X'\left(\sum_{i=1}^{\infty} \rho^i W^i\right)X\beta \qquad [13]$$

From equation 11, or equation 12, the expected value of $\hat{\beta}$ can be obtained, and from equation 13 the magnitude of the bias. Further, this bias will increase with ρ.

The true sample variability of $\hat{\beta}$ is readily obtained as is the OLS report of the standard error. From equations 10 and 11,

$$\hat{\beta} - E\hat{\beta} = (X'X)^{-1}X'A^{-1}\epsilon,$$

and so

$$(\hat{\beta} - E\hat{\beta})(\hat{\beta} - E\hat{\beta})' = (X'X)X'A^{-1}\epsilon\epsilon'A^{-1'}X(X'X)^{-1}$$

Taking expectations, we have

$$V\hat{\beta} = \omega(X'X)^{-1}X'A^{-1}A^{-1'}X(X'X)^{-1} \qquad [14]$$

In equation 14 if $\rho = 0$, then $V\hat{\beta} = \omega(X'X)^{-1}$, which is the OLS result. Clearly, as ρ departs from zero it is likely that the true $V\hat{\beta}$ is not reported correctly by the OLS results. Writing $V\hat{\beta} = \omega(X'X)^{-1}Q$, the behavior of Q provides some indication of the extent to which the OLS report departs from the true sample variability of the coefficient estimates. For OLS, $Q = I$. Although equation 14 is markedly nonlinear in ρ, it can still be used to show the departure of $\omega(X'X)^{-1}$ from $V\hat{\beta}$ (if ω is known).

However, OLS does not report $\omega(X'X)^{-1}$ as it first has to estimate ω. Therein lies another problem with OLS. Consider $\epsilon'\epsilon$ in equation 8. By definition,

$$\begin{aligned}\hat{\epsilon} &= y - X\beta \\ &= y - X(X'X)^{-1}X'y \\ &= [I - X'(X'X)^{-1}X']\, y = My \qquad [15]\end{aligned}$$

If the OLS specification held, then $\hat{\epsilon} = M\epsilon$ and $E\,\hat{\epsilon}'\hat{\epsilon} = \omega\mathrm{tr}\,(M)$ giving $E\hat{\epsilon}'\hat{\epsilon} = (N - K - 1)\omega$ (Goldberger, 1964). Which means ω as given in equation 8 is an unbiased estimator of ω. But for network autocorrelation, from equation 15, we have

$$\hat{\epsilon} = M(A^{-1}X\beta + A^{-1}\epsilon)$$

and

$$E\,\hat{\epsilon}'\hat{\epsilon} = E[\epsilon A^{-1'}M'MA^{-1}\epsilon] + \beta'X'A^{-1'}MA^{-1}X\beta$$

which reduces to

$$= \omega \mathrm{tr}(A^{-1'}MA^{-1}) + \beta'X'A^{-1}MA^{-1}X\beta \qquad [16]$$

If $\rho = 0$, $A^{-1} = I$ and equation 16 reduces to $E\hat{\epsilon}'\hat{\epsilon} = (N - K - 1)\omega$ and $E\hat{\omega} = \omega$ (which are the OLS results). Thus we would expect OLS to be in error when it reports a coefficient estimate and its standard error.

SIMULATION RESULTS[8]

THE FULL MODEL

The full model is simply one in which none of the generating β parameters is set to zero when the data sets are created. We consider first the case where there is no network autocorrelation and then introduce network autocorrelation.

No Network Autocorrelation

For $\rho = 0$, Table 1 gives the results when the disturbance term is set to 121. Considering β_1, the mean parameter estimates are all close to the generating value of 0.3. On rounding, all methods return mean parameter estimates equal to the generating value. As far as the standard deviation of the estimate is concerned, MLE and OLS are close to each other and each is considerably less variable than QAD. This makes sense, for the inclusion of Wy as a regressor on the left-hand side of equation 2 amounts to inclusion of an irrelevant regressor when $\rho = 0$, which leads to lower efficiency (Rao and Miller, 1971). For both QAD and OLS the mean reported standard error is less than the actual standard deviation of the estimate, while for the MLE the mean reported

TABLE 1
Parameter Estimates and Performance Measures for Runs with No Network Autocorrelation: ω = 121

Generating Parameter		Mean Parameter Estimate (MPE)	Standard Deviation of Estimate (SDE)	Mean Reported Standard Error (MRSE)	Range of Estimate Min	Range of Estimate Max
β_1=0.3	MLE	0.3000*	0.0404	0.0469	0.1872	0.3813
	QAD	0.3082	0.0679	0.0534	0.1651	0.5298
	OLS	0.2984	0.0394	0.0367	0.1874	0.3755
β_2=-0.3	MLE	-0.3026	0.0568	0.0547	-0.4486	-0.1012
	QAD	-0.3002	0.0576	0.0557	-0.4505	-0.0962
	OLS	-0.3026	0.0569	0.0556	-0.4484	-0.0996
β_0=10.0	MLE	10.20	2.80	2.63	3.70	16.35
	QAD	10.22	2.97	2.72	3.31	19.14
	OLS	10.18	2.77	2.58	3.99	16.27
ω=121	MLE	116.63	19.52	20.62	76.84	174.96
	QAD	119.91	20.57	21.71	74.87	183.88
	OLS	120.60	20.50	21.66	78.88	181.12
ρ=0	MLE	-0.007	0.027	0.152	-0.031	0.164
	QAD	-0.046	0.222	0.196	-0.639	0.439

*Actually 0.29997.

standard error is above the actual standard deviation.[9] Finally, QAD has a wider range for its estimates consistent with its having a higher standard deviation. Focusing attention on β_2, we see far less variation across the estimation methods. All report a mean parameter estimate equal to the generating value and the standard deviations of the estimate are close. Further, the mean reported standard error is close for all methods and all are close to their corresponding standard deviation of the estimate; the ranges are virtually identical. This suggests, and a subsequent simulation bears this out, that there will be greater variability in terms of performance when parameter β_1 is considered compared to that of β_2.[10]

For the intercept, the mean parameter estimates for all methods are close to each other, and the generating value of 10, while the standard deviation of the estimates for MLE and OLS are close. Both are slightly lower than the standard deviation of the estimate for QAD. All three methods underreport the standard deviation of the estimate for $\rho = 0$. For ω, the MLE estimate is below the QAD and OLS estimates, which are close to each other. For the standard deviation of the estimate, MLE is less than QAD and OLS (which are again close). For all estimation methods the mean reported standard error is above the actual standard deviation.

Finally, we consider ρ. Obviously, for $\rho = 0$, OLS is preferable as its expected value is 0, its bias is 0, and its standard error is 0. The mean MLE estimate is closer to the actual generating parameter of 0 than the QAD estimate and it has much smaller variability across the estimates. Indeed, the QAD standard deviation seems alarmingly large. In terms of the mean reported standard error, we find the MLE report to be considerably above its actual standard deviation, whereas the QAD mean reported standard error is lower than the actual standard deviation. The range of the MLE estimate is smaller than the range of the QAD estimate. In fact, a range from –0.6 to +0.4, although containing the generating value of 0, is large (and reflected in the high actual standard deviation).

As far as $\hat{\beta}$ is concerned, for $\rho = 0$, all three methods perform well, as would be expected. The biases are low or nonexistent and the sample variances are comparable. If anything, QAD is more variable than either OLS or MLE. For estimating $\rho(=0)$ MLE is closer to zero than QAD. It is much less variable across samples, with QAD appearing inadequate. Further, the mean MLE report of sample variability is much higher than the actual sample variability (by a factor of around 6), whereas the reverse is true for QAD. This pattern persists through the simulations, making inference about $\rho = 0$ conservative for MLE and nonconservative for QAD.

The simulation results for other values of ω (but with ρ still set to 0) are very simply summarized. The mean parameter estimate

TABLE 2
Estimates of β_1 in the Full Model for ω = 121 and $0.1 \leqslant \rho \leqslant 0.9$

Estimation Method	Values of ρ	Mean	Bias	SDE	Root Mean Square Error (RMSE)	MRSE	Range of Estimates Min	Max
MLE	.1	0.3002	0.0002	0.0405	0.0405	0.0482	0.1873	0.3818
	.3	0.3009	0.0009	0.0406	0.0406	0.0506	0.1876	0.3829
	.5	0.3016	0.0016	0.0408	0.0408	0.0532	0.1880	0.3841
	.7	0.3027	0.0027	0.0408	0.0409	0.0546	0.1884	0.3853
	.9	0.3044	0.0044	0.0408	0.0410	0.0551	0.1889	0.3864
QAD	.1	0.2992	0.0008	0.0690	0.0690	0.0552	0.1567	0.5270
	.3	0.2791	-0.0209	0.0598	0.0633	0.0585	0.1409	0.5100
	.5	0.2581	-0.0419	0.0687	0.0805	0.0614	0.1204	0.4770
	.7	0.2414	-0.0586	0.0665	0.0886	0.0637	0.0989	0.4325
	.9	0.2464	-0.0536	0.0638	0.0833	0.0632	0.0999	0.3842
OLS	.1	0.3205	0.0205	0.0424	0.0471	0.0368	0.1986	0.4027
	.3	0.3827	0.0827	0.0512	0.0972	0.0383	0.2301	0.4785
	.5	0.4893	0.1893	0.0666	0.2007	0.0429	0.2833	0.6140
	.7	0.7136	0.4136	0.1001	0.4255	0.0564	0.3908	0.9160
	.9	1.5422	1.2422	0.2326	1.2638	0.1244	0.7469	2.0802

is always on target for the regression coefficients. The standard deviation of the estimate increases with ω as does the mean reported standard error.[11] Of course, the range of the estimate tends to broaden with increasing ω. All of the above is as expected. The variability in the performance in the estimation method is greater for β_1 than it is for β_2. The relative magnitudes of the mean reported standard errors remain the same across estimation methods.

In summary, the behavior of the parameter estimates for each of the estimation strategies is straightforward and forms a point of departure for the introduction of different degrees of network autocorrelation.

Network Autocorrelation

With ω = 121, we consider the full model with ρ ranging from 0.1 to 0.9. The results are displayed in Table 2. The standard deviation of the estimate, SDE, is the actual standard deviation of the estimates (for each parameter) over all 100 runs. This measures the actual variability of an estimator. For each run, a value for the estimate of sample variability is computed. The mean value of these is the mean reported standard error (MRSE) and is taken to characterize the reports of sample variability provided by an estimator. As is expected theoretically, the estimates of β_1 are unbiased with MLE. The empirical approximation to the amount of bias never exceeds 0.0044, which seems trivially small.[12] The direct measures of the sample variability (SDE) remain virtually fixed throughout the range of ρ at around 0.04. Consistent with the results for $\rho = 0$, the mean reported standard error from this procedure exceeds the standard deviation of the estimates.[13] The MLE estimates range from around 0.19 to around 0.39.

For the QAD method of estimation there appears to be a downward bias in the estimate of β_1. Although this bias is slight for small values of ρ, it does increase with increasing ρ. By $\rho = 0.5$, this downward bias is around 14% relative to the true value and reaches around 18% for ρ equal to 0.9. The increase in the magnitude of this bias appears much more consequential than the slight increase of the bias under the MLE method. When we consider the standard deviation of the estimates there is a slight decrease in the variability of these estimates as ρ increases.[14] Although the root mean square error (RMSE) provides a composite estimate of bias and sample variability, it is not necessary to compute it for a comparison between the MLE and the QAD methods. For each value of ρ, the standard deviation is higher for the QAD method, as is the value of the bias. So the RMSE for MLE is always lower than for QAD. For QAD, the mean reported standard error indicates that it always underreports the true variability of its estimates. However, as SDE tends to decrease and MRSE tends to increase as ρ increases, the magnitude of the underreporting decreases with ρ. To the extent that t-ratios are constructed with an estimate of the sample variability in the denominator, it is reasonable to surmise that a bias in the estimation of standard error is in the direction of finding more significant parameters than may be the case.

Finally, the range of the parameter estimates under the QAD method is considerably larger than under the MLE method (which is consistent with the higher standard deviation).

The OLS data, comparable to those for MLE and QAD, are found in the third panel of Table 2. As a method it fares badly. It is biased upward in its estimation of β_1, with the bias getting larger as ρ increases. For high ρ the bias is absurdly high. It is noteworthy that for $\rho = 0.5$, the range of the estimate barely includes the generating value of 0.3, whereas for $\rho = 0.7$ it no longer includes the generating value. The direct estimate of the sample variability, SDE, increases continually with ρ. For $\rho = 0.1$, the variability of the OLS estimate is comparable to that of the MLE method. Although the MLE variability measure remains constant, that for the OLS method increases rapidly. For small values of ρ, the OLS sample variability is lower than that for the QAD method and remains so until around 0.5. As it is more likely to have values of network autocorrelation ranging between 0 and 0.5, it is worth noting that OLS estimation method for this parameter does have lower variability over this part of the range of ρ. In terms of RMSE measure, however, OLS is preferable only for $\rho = 0.1$; thereafter the RMSE for QAD is lower than for OLS. The mean reported standard error for the OLS method is always less than the standard deviation of the estimates, suggesting that when the downward bias of the reported standard error is coupled to upward bias in the estimate of β_1, there will be many instances of a parameter being deemed significantly different than 0 when its value is 0.

Of course, we would not need the simulation for much of the third panel of Table 2 if we were to focus on OLS alone. The analytical results presented earlier can be used to give the performance of OLS. Use of equation 12 gives $E\hat{\beta}$ and use of equation 13 gives the bias of $\hat{\beta}$; these are shown in the first two panels of Table 3. The close correspondence between the theoretical and empirical values indicates that the data were generated correctly. The third panel of Table 3 gives the expected standard errors from use of equation 14 with ω taken as 121, and the fourth panel gives the extent to which OLS underreports the

TABLE 3
Bias and Sample Variability of OLS in the Presence of Network Autocorrelation: ω = 121

		Value of rho				
Property	Parameter	.1	.3	.5	.7	.9
Expected Value	β_1	0.3221	0.3841	0.4902	0.7130	1.532
	β_2	-0.3015	-0.3102	-0.3304	-0.3760	-0.5080
	β_0	10.46	11.99	15.19	23.80	76.97
Bias	β_1	0.0221	0.0841	0.1902	0.4130	1.223
	β_2	-0.0015	-0.0102	-0.0304	-0.0760	-0.2080
	β_0	0.46	1.99	5.19	13.80	66.97
Theoretical Standard Error (Known ω)	β_1	0.040	0.049	0.065	0.099	0.234
	β_2	0.056	0.058	0.064	0.077	0.118
	β_0	2.753	3.262	4.236	6.53	17.41
Ratio to OLS Report	β_1	1.084	1.323	1.754	2.681	6.322
	β_2	1.005	1.043	1.140	1.375	2.102
	β_0	1.063	1.259	1.635	2.521	6.720

standard error (if ω is taken as known). This panel clearly shows a persistent tendency for OLS to underreport the true sample variability for each of the regression parameters with this regime.

In summary, the performance of MLE is the best of the three estimation methods, with that of OLS being the worst. We will present a comparison of MLE and QAD in terms of inference decisions after a discussion of the bias and sample variability issues for the remaining parameters of the model.

From examination of the results in Table 1 it is reasonable to suspect that the estimates of β_2 are less sensitive to changes in the values of the network autocorrelation parameter. This expectation is borne out by the results in Table 4. Given the very close nature of the estimation results for $\rho = 0$, it is not surprising to see very similar results for $\rho = 0.1$ in Table 4. The estimation methods are indeed close. Throughout the range of ρ the mean estimate of

TABLE 4
Estimates of β_2 in the Full Model for ω = 121 and $0.1 \leqslant \rho \leqslant 0.9$

Estimation Method	Values of ρ	Mean	Bias	SDE	RMSE	MRSE	Range of Estimates Min	Max
MLE	.1	-0.3026	-0.0026	0.0568	0.0569	0.0548	-0.4488	-0.1011
	.3	-0.3028	-0.0028	0.0569	0.0570	0.0549	-0.4492	-0.1008
	.5	-0.3030	-0.0030	0.0569	0.0570	0.0551	-0.4495	-0.1004
	.7	-0.3033	-0.0033	0.0567	0.0568	0.0554	-0.4497	-0.1000
	.9	-0.3035	-0.0055	0.0570	0.0572	0.0553	-0.4496	-0.09958
QAD	.1	-0.3003	-0.0003	0.0575	0.0575	0.0558	-0.4524	-0.0972
	.3	-0.2985	0.0015	0.0575	0.0575	0.0559	-0.4542	-0.0990
	.5	-0.2952	0.0048	0.0574	0.0576	0.0559	-0.4589	-0.1003
	.7	-0.2918	0.0082	0.0571	0.0577	0.0559	-0.528	-0.1008
	.9	-0.5038	0.0062	0.0566	0.0569	0.0559	-0.4541	-0.1000
OLS	.1	-0.3038	-0.0038	0.0572	0.0570	0.0558	-0.4503	-0.0978
	.3	-0.3118	-0.0118	0.0594	0.0606	0.0581	-0.4615	-0.0938
	.5	-0.3312	-0.0312	0.0652	0.0723	0.0651	-0.4882	-0.0890
	.7	-0.3752	-0.0752	0.0792	0.1092	0.0854	-0.5476	-0.0823
	.9	-0.5029	-0.2029	0.1222	0.2269	0.1886	-0.7741	-0.0732

the MLE of β_2 remains fixed and close to the true parameter. The mean value of β_2 for the QAD method remains close to the true parameter. Indeed, for lower values of ρ the QAD method is closer to the generating parameter than the MLE method, although the differences are very small.[15] The standard deviation of the estimates remain fixed for both MLE and QAD at around 0.057. Further, the mean reported standard error for both the MLE and QAD methods are close. In terms of the criteria reported in Table 4, no real difference exists between MLE and QAD for the estimation of this parameter.

For the bias criterion, the standard error of the estimate criterion, and their RMSE composite, OLS performed progressively worse, just as it did with the estimation of β_1. There is a continual increasing downward bias (away from 0) in the estimation of β_2 and an increase in the standard deviation of the estimates. However, the problem of underreporting via the mean standard error is less prevalent for estimating this parameter. Although for small

TABLE 5
Estimates of β_0 (intercept) in the Full Model for ω = 121 and $0.1 \leqslant \rho \leqslant 0.9$

Estimation Method	Values of ρ	Mean	Bias	SDE	RMSE	MRSE	Range of Estimates Min	Max
MLE	.1	10.21	0.21	2.80	2.81	2.65	3.69	16.37
	.3	10.23	0.23	2.81	2.82	2.71	3.68	16.40
	.5	10.26	0.26	2.81	2.82	2.80	3.69	16.48
	.7	10.31	0.31	2.82	2.84	2.97	3.76	16.49
	.9	10.49	0.49	2.82	2.83	3.51	4.05	16.63
QAD	.1	10.05	0.05	2.92	2.92	2.75	3.59	19.36
	.3	9.59	-0.41	2.82	2.85	2.83	3.93	19.62
	.5	8.97	-1.03	2.75	2.94	2.94	3.44	19.63
	.7	8.20	-1.80	2.76	3.30	3.16	3.01	19.50
	.9	7.42	-2.58	3.40	4.27	3.95	2.38	20.08
OLS	.1	10.63	0.63	2.97	3.04	2.58	4.13	17.12
	.3	12.12	2.12	3.58	4.16	2.69	3.59	19.80
	.5	15.26	5.26	4.74	7.08	3.01	3.17	25.47
	.7	23.70	13.70	7.43	15.58	3.95	3.84	39.97
	.9	76.02	66.02	20.33	69.08	8.74	22.56	120.87

ρ there is a tendency to underreport, this is no longer true for ρ = 0.5 and, for higher ρ, the mean reported standard error is above the actual standard deviation. The bias in β_2 for OLS in the direction away from 0 is not coupled with a bias that tends to minimize the report of sample variability. For smaller values of ρ, MRSE is smaller than SDE but the differences seem minor. For the higher values of ρ, MRSE does not underreport SDE, but the bias increases quickly. The range of the estimates for β_2 tends to be greater for OLS than for MLE or QAD, but it does contain the generating parameter at all times, unlike the case for β_1.

Even though the magnitude of the intercept term tends to be of lower substantive importance than the values of the other parameters in most estimated linear models, it is worth reporting, nevertheless. The results are shown in Table 5. The MLE method targets on the parameter value with a slight upward bias. The standard deviation of the estimates hardly varies with increases in the network autocorrelation parameter ρ. For low values of ρ

there tends to be a slight downward bias in the mean reported standard error relative to the actual standard deviation for lower values of ρ, rough parity for $\rho = 0.5$, and a slight upward bias for the higher values of ρ. For $\rho = 0.1$, the mean $\hat{\beta}_0$ for QAD is close to the generating value of 10, but declines as ρ increases. Thus, for the higher values of ρ, there is a downward bias. For the range $0.1 \leq \rho \leq 0.7$ the standard deviation of the estimates declines slightly but is close to the MLE figures. For $\rho = 0.9$ it jumps upward (but only at this extreme). The mean reported standard error for QAD is above the actual standard deviation of the estimates for $\rho \geq 0.3$. For all ρ reported in Table 5, the range of the estimate is higher for QAD than for MLE.

Although the OLS estimate is close to the generating value of 10 for $\rho \leq 0.1$ it rapidly increases with ρ so that there is a marked upward bias throughout the range of $\rho > 0.1$. Even for $\rho = 0.3$ the bias is considerable (21%). The standard deviation of the estimates also rises continually with the increases in ρ. The mean reported standard error also rises but it does so much less rapidly, providing an underreporting of the true variability. However, only with the extreme value of $\rho = 0.9$ does the range of the OLS estimator not contain the generating parameter.

For the estimation of the intercept, the simulation results point to the superiority of the MLE method and the inferiority of the OLS method. However, the behavior of QAD is comparable to that of MLE (see RMSE figures), with the exception of the high, and empirically unlikely, values of the network autocorrelation parameter.

Table 6 reports the estimation results in the simulation for the estimate of the variance in the disturbance term. For all values of ρ, the mean of $\hat{\omega}$ under MLE has approximately a 4% downward bias relative to the generating value of 121. Although the mean value of the estimate of $\hat{\omega}$ is also biased downward for QAD, the bias is smaller. The standard deviation of the estimates show that MLE is slightly less variable than QAD but the differences between them are small. For the mean reported standard error, a tendency exists to report a value higher than the actual standard deviation of the estimate for both methods and to do so by a

TABLE 6
Estimates of ω in the Full Model Where ω = 121 for $0.1 \leqslant \rho \leqslant 0.9$

Estimation Method	Values of ρ	Mean	Bias	SDE	RMSE	MRSE	Range of Estimates Min	Max
MLE	.1	116.67	-4.33	19.53	20.00	20.64	76.90	175.14
	.3	116.68	-4.32	19.54	20.00	20.76	77.01	175.53
	.5	116.90	-4.10	19.57	19.99	20.96	77.14	175.93
	.7	117.05	-3.95	19.61	20.00	21.17	77.24	176.29
	.9	117.21	-3.79	19.66	20.02	21.19	77.22	175.25
QAD	.1	120.08	-0.92	20.54	20.56	21.74	74.13	183.23
	.3	119.82	-1.18	20.44	20.47	21.70	72.80	180.93
	.5	119.00	-2.00	20.23	20.33	21.55	71.96	180.10
	.7	118.15	-2.85	19.91	20.11	21.39	72.27	178.83
	.9	118.62	-2.38	19.90	20.04	21.48	75.16	179.40
OLS	.1	121.33	-0.33	20.86	20.86	21.79	82.55	138.50
	.3	131.53	10.53	24.11	26.31	23.62	84.93	217.16
	.5	165.38	44.38	34.70	56.33	29.70	96.28	303.91
	.7	286.36	165.36	74.62	181.42	51.14	131.19	574.75
	.9	1421.6	1300.6	498.41	1392.8	255.33	301.74	2934.0

comparable amount. Further, the ranges of the estimates are comparable for both MLE and QAD. In effect, the performance of QAD is virtually indistinguishable from that of MLE, although the MLE RMSE is always smaller than the RMSE for QAD.

OLS again fares badly. Here, an upward bias increases dramatically, with ω reaching absurdly high levels for high ρ. The standard deviation of the estimates also increases with ρ, as does the mean reported standard error, although, for the most part, it underreports the magnitude of the variability of the OLS estimator of ω. The RMSE rapidly increases with ρ.

Finally, for the full model, we consider the estimates of ρ. These results are shown in Table 7. Obviously, the implicit OLS estimate of 0 for the network autocorrelation parameter renders it biased for all nonzero values of ρ and this bias increases throughout the range of ρ. Equally obvious is that no sample variability of the estimate exists and an RMSE measure is composed of the bias only. However, this bias component for

TABLE 7
Estimates of ρ in the Full Model for ω = 121 and $0.1 \leqslant \rho \leqslant 0.9$

Estimation Method	Values of ρ	Mean	Bias	SDE	MRSE	Range of Estimates Min	Max
MLE	.1	0.0920	-0.0080	0.0251	0.1454	0.0700	0.2517
	.3	0.2914	-0.0086	0.0270	0.1279	0.2780	0.4219
	.5	0.4917	-0.0083	0.0152	0.1038	0.478	0.5868
	.7	0.6930	-0.0070	0.0088	0.0724	0.6251	0.7479
	.9	0.8957	-0.0043	0.0027	0.0311	0.8925	0.9112
QAD	.1	0.0996	-0.0004	0.2087	0.1867	-0.4666	0.5406
	.3	0.3729	0.0720	0.1713	0.1627	-0.1129	0.7127
	.5	0.6083	0.1083	0.1260	0.1300	0.2236	0.8430
	.7	0.7978	0.0978	0.0790	0.0890	0.5517	0.9394
	.9	0.9375	0.0375	0.0326	0.0384	0.8489	0.9660

OLS will completely outweigh both the bias and standard deviation components of MLE and QAD, rendering OLS a poor estimator of ρ. In terms of the mean value of ρ over the simulations, both MLE and QAD work well. Throughout most of the range the slight downward bias of the MLE is smaller in magnitude than the slight upward bias of QAD. However, when we consider the sample variability of the estimates, measured by SDE, the MLE method is much preferable to that of QAD. Although for both MLE and QAD the standard deviation measure drops, surprisingly, with increases in the value of ρ, the variability is higher for QAD relative to MLE. (For $\rho = 0.1$ it is around 8 times higher and the discrepancy has a factor of 14 for $\rho = 0.9$.) The RMSE criterion shown in Table 7 indicates the marked supremacy of the MLE method over the QAD as far as estimating the network autocorrelation parameter is concerned. The two methods are less distinguishable when the MRSE is concerned. Both overreport relative to the actual standard deviation, but the MLE method does so to a far greater extent. The marked upward bias in the estimate of the sample variability renders the MLE approach more conservative than is perhaps warranted. For $\rho = 0.1$ it overreports by a factor of 6 and by a factor of 12 for $\rho = 0.9$. In terms of the range of the estimates both the MLE and the QAD methods bracket the true generating parameter. However, consistent with the results on the standard deviation of the estimates, the MLE bracket is much tighter than that of QAD. In fact, for the lower values of ρ, the QAD range seems distressingly large.

The results for the runs with ω set to 25, 49, and 81 are reported in two ways: (1) with each batch considered separately and (2) all batches considered together. As is expected from the mathematical argument presented earlier, the bias for OLS is completely unaffected by the variance of the disturbance term: The same pattern of increasing bias with increasing values of ρ for OLS is found in each batch. The standard deviation of the estimates increases as ω increases in the simulation runs for each estimation method, but the pattern and relative behavior of these estimates is much the same for the other values of ω as for the case of $\omega = 121$. The mean reported standard error also increases with ω but—in relation to the standard deviations of the estimate—the same relative pattern as for $\omega = 121$ is found. In short, the relative behavior of MLE, QAD, and OLS is unchanged with changes of ω, but the numerical differences are magnified with increasing ω.

When all the runs, with $\rho \neq 0$, are considered together, some crude quantitative descriptions can be established. The first panel of Table 2 provides the right-most swarm of data points in Figure 1, where the MLE mean reported standard error has been plotted against the MLE standard deviation of the estimate for β_1. The scatter plot is, in large measure, driven by the four swarms corresponding to the four values of the disturbance term variance (ω). Within each vertical swarm the values are ordered according to ρ. If the mean reported standard error is regressed on the standard deviation of the estimates and ρ we have:[16]

$$\underset{}{\text{MRSE}} = \underset{(0.001)}{0.007} + \underset{(0.033)}{1.045}\ \text{SDE} + \underset{(0.001)}{0.006}\,\rho \qquad R^2 = 0.98 \qquad [17]$$

From equation 17, for which fitted values are close to the actual MRSE, we see MRSE is always above SDE and the gap grows with increasing ρ.

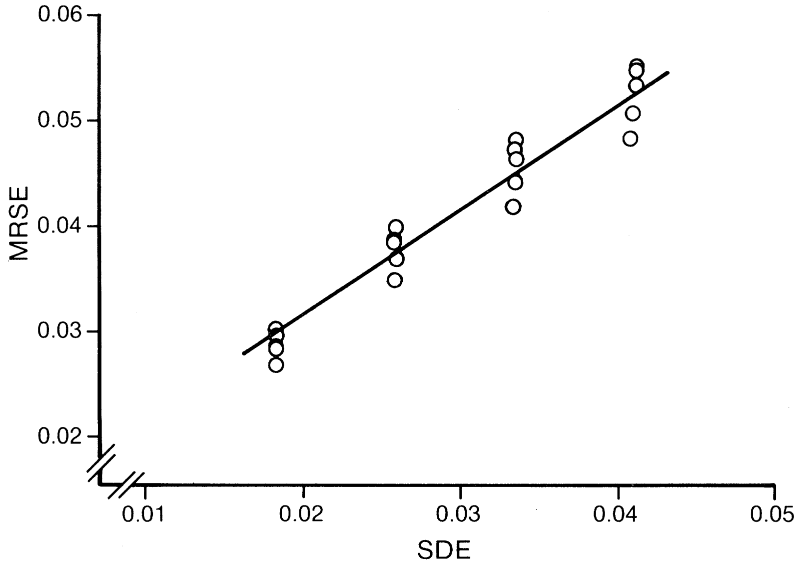

Figure 1: Plot of MRSE against SDE, Full Model for β_1 with ω = 121

For QAD, the plot of MRSE against SDE has the same general form but the numerical values for the regression[17] differ:

$$\underset{}{\text{MRSE}} = \underset{(0.042)}{0.841}\ \text{SDE} + \underset{(0.002)}{0.009}\ \rho \qquad R^2 = 0.96 \qquad [18]$$

This numerical summary, for all ω, is reflected in each ω batch. For the second panel of Table 2, we see that MRSE is smaller than SDE and the gap narrows with increases in ρ.

For OLS no distinct swarms of points are associated with different values of ω as the ranges of MRSE and of SDE for different ω overlap. Nor does ρ, treated as a regressor, have an appreciable effect on MRSE. The following is a crude curvilinear[18] description:

$$\text{MRSE} = \underset{(0.007)}{0.006} + \underset{(0.165)}{0.966}\ \text{SDE} - \underset{(0.688)}{1.673}\ (\text{SDE})^2 \qquad R^2 = 0.93 \qquad [19]$$

Taking partial derivatives, $\partial MRSE/\partial SDE = 0.966 - 1.673$ SDE so that the gap between the OLS mean reported standard error and the true standard deviation of the estimates narrows as the latter increases (which is driven by ω and ρ).

In this fashion a crude quantitative statement of the relation between the mean reported standard error and the true standard deviation of the estimates can be found for each estimation method and each estimated parameter.[19] This quantitative statement, together with the statement that the same general qualitative behavior as that found in Tables 2 and 4 through 7 is found for the runs with the different ω, provides a sufficient summary of the results, other than inference, of all the simulations featuring the full model.

Inference Decisions with Network Autocorrelation

We consider inference and deal with two kinds of null hypothesis: (1) H_0: $\theta = 0$ and (2) H_0: $\theta = \theta_{gen}$. Throughout this discussion we use $\alpha = 0.01$ as a significance level.

From the viewpoint of testing whether or not a variable is significant in a linear equation—that is, the first kind of null hypothesis—the results of the simulation for the full model are of little interest. Consider any parameter other than ρ. As can be seen by the small nature of the bias for MLE and QAD, and—relative to the parameter estimates—the small measure of the mean reported standard error of the estimates, all runs show any parameter to be significantly different from 0. For OLS, the biases are away from zero and, as the OLS report of the standard errors are biased downward, it, too, reports significant β_i. But it is inappropriate to draw any conclusions, as the parameter values were all chosen to be large enough so as to ensure that they would be detected as nonzero. Issues of inference when one of the generating parameters may really be 0 becomes the crucial area in which to explore questions of inference. We deal with this shortly.

A second and, herein, more important question has to do with the actual value of the estimated parameter. Given the appropriate specification of a linear model, the values of the estimated parameters are important. The extent to which the obtained numerical values can be taken seriously depends on how close to the true parameters we can expect them to be. If we know they are likely to be reasonably close to the true value, our confidence in them increases. This can be formulated as an inferential question. If we specify as a null hypothesis the actual parameters used to generate the data, we should not be able to reject such a null hypothesis in most instances. To the extent that an estimation method, and its coupled inference procedure, leads to the rejection of the second kind of null hypothesis, the whole issue of the correct estimation of the value of the parameter is called into question. Stated differently, if an estimation procedure returns estimates of parameters (together with their standard errors) that lead frequently to the rejection of the null hypothesis stating that the parameter has the generating value, we can have little confidence in the accuracy of the estimation procedure. With this issue in focus we consider Table 8.

This table shows the count of the correct nonrejection of the null hypothesis stating that parameter values are those used in the generating procedure. This is shown for all values of ρ in the simulation range and all values of disturbance term variance. We seek to establish, for each estimation procedure, if the quality of the inference decisions declines as ρ ranges from 0 to 0.9 and as ω ranges from 25 to 121. The broad result is that MLE and QAD do not appear to decline in inference performance as ρ increases, whereas OLS does (in a spectacular fashion). Second, both MLE and QAD do not decline in quality of inference decisions as ω increases and, strangely, the quality of the OLS decision appears to improve as the disturbance term variance increases. However, this improvement is not sufficient to offset the severe degradation in the quality of the inference decision for OLS as ρ increases.

We consider first the coefficient β_1. When ρ is 0 there is not much difference between MLE and OLS—or at least the differences between them may be attributable to random sources. MLE, according to this table, will be in error 1 in 100 times, whereas OLS will be in error 5 times in 100 trials. QAD may be slightly inferior to either of the other two methods insofar as it

TABLE 8
Count of Correct Nonrejection of $H_o : \theta = \theta_{gen}$ for Full Model (.3, - .3, 10)

		ω=25			ω=49			ω=81			ω=121		
		MLE	QAD	OLS	MLE	QAD	OLS	MLE	QAD	OLS	MLE	QAD	OLS
β_1	0	99	92	95	99	90	95	98	89	95	98	87	95
	0.1	99	93	67	99	93	81	98	89	81	98	85	86
	0.3	99	94	2	99	90	17	99	89	29	98	88	47
	0.5	99	92	0	99	88	0	99	87	2	98	85	4
	0.7	100	92	0	99	89	0	99	84	0	99	80	0
	0.9	99	92	0	99	90	0	99	89	0	99	86	0
β_2	0	93	94	93	93	93	93	93	93	93	93	93	93
	0.1	93	94	93	93	94	93	93	94	93	93	93	94
	0.3	93	94	95	93	94	94	93	94	93	93	94	95
	0.5	93	95	93	93	95	93	93	95	95	93	95	95
	0.7	94	95	91	93	95	90	93	95	91	93	95	92
	0.9	93	94	100	93	95	97	93	95	96	93	95	96
β_0	0	93	92	93	93	93	93	93	95	93	93	94	93
	0.1	93	94	87	93	94	89	93	93	89	93	93	90
	0.3	93	94	63	93	94	71	93	94	74	93	94	81
	0.5	95	95	19	95	95	33	94	93	49	94	93	54
	0.7	98	94	1	94	95	5	96	94	14	94	93	22
	0.9	100	94	0	99	95	0	99	95	0	98	95	1
ω	0	93	94	94	93	94	94	93	94	94	93	93	94
	0.1	93	95	95	93	94	95	93	94	95	93	94	94
	0.3	93	95	92	93	95	95	94	96	97	94	96	97
	0.5	93	96	14	94	96	50	94	96	73	94	96	98
	0.7	93	96	0	94	96	1	94	96	2	94	96	7
	0.9	93	95	0	94	95	0	94	95	0	94	95	0

leads to the incorrect rejection of the null hypothesis around 10 times out of 100. There is also mild evidence (for this parameter only) that the quality of the decision making for QAD may decrease as ω increases. This is reasonable for—with $\rho = 0$—the QAD essentially includes an irrelevant variable that raises slightly the estimate of the standard error of the coefficient estimates. As ρ increases, the error rate of the MLE remains around 1 in 100 for all values of ω. As ρ increases, the error rate for QAD is around 8 times in 100 for $\omega = 25$, around 10 in 100 for $\omega = 49$, around 11 times in 100 for $\omega = 81$, and around 12 in 100 for $\omega = 121$. If anything, there is a slight increase in the error rate for QAD inference for $\rho = 0.5$ and 0.7, although these may themselves be random departures from a fixed pattern. For OLS, however, the rate of correct inferences drops sharply as ρ increases. Already by $\rho = 0.1$, for $\omega = 25$ there are only 67 correct inferences out of 100. By the time ρ reaches 0.3 the number of correct inferences is only 2 out of 100 and for the higher values of ρ a correct inference never exists concerning the null hypothesis of β_1 being set to its generating value. There are similar declines in the rate of error making for OLS for the other three values of ω, although the rate of decline is not as sharp, as ρ increases.

For the estimates of β_2 we predicted much more stability on the basis of the earlier results; this is borne out in the second panel of Table 8. Essentially, the error rate for all methods, over all values of ρ and ω, is around 7 times in 100 for H_0: $\beta_2 = -0.3$. The stability of this error rate across all values of ρ and ω is striking and is in stunning contrast to the results for the other slope parameter.

As far as inference affecting the intercept is concerned, the pattern for MLE and QAD is similar to that of β_1, with perhaps a slightly higher rate of error making for MLE and a slightly lower rate for QAD. For both MLE and QAD, the inference decision is wrong about 7 times in 100 under all combinations of ρ and ω displayed in the table. For OLS, however, the numbers of correct inferences concerning β_0 decline markedly as ρ increases, for all values of ω. For higher values of ω, there tends to be a slower increase in the rate of incorrect inference as ρ increases but, for all ω, they are alarmingly high. Similar results hold true for the estimation of ω, with MLE showing 7 wrong inference decisions out of 100 for all ρ and ω, with QAD showing slightly fewer incorrect decisions. Again a decline in the number of correct inferences made under OLS occurs. However, the decline is much less rapid than is the case for β_1 and β_2. For low values of ρ, through 0.3, OLS is comparable to both MLE and QAD. Thereafter the rate of correct inference drops sharply with increases in ρ, except for $\omega = 121$, for which the sharp drop is delayed until after $\rho = 0.5$.

In summary, the implication of the results displayed in Table 8 is striking. At least for the combination of parameters considered in these simulations, OLS is inadequate in rendering correct decisions concerning the actual generating values, whereas MLE and QAD perform much better and with comparable quality.[20] Of course, social science researchers are seldom in a position to test null hypotheses positing nonzero values for parameters. We know already from the mathematical argument presented above that the OLS estimators would be biased, and analytically, we could provide statements concerning the reported standard errors. But those separate results could not be combined readily. The simulations show in this context that the OLS values tend to be not only different from the generating values, but also sufficiently far away from the true values that OLS reports them as significantly different from the generating values. An estimation procedure focusing on the wrong target, having large sample variability, and leading to many wrong inferences merits no confidence whatsoever in the accuracy of the estimates it provides.

MODELS WITH ONE ZERO SLOPE COEFFICIENT

Parameter Estimation

Loftin and Ward (1983) present an example in which OLS returns a significant coefficient for a specific variable and an estimation method incorporating network autocorrelation does not. In our results for the full model we saw that the OLS estimate of β_1 was biased upward, whereas the OLS estimate of the standard error was biased downward. Together the two biases may lead to mistaken inference concerning the presence of β_1 in an equation. A reasonable speculation, then, is that such a mechanism was at work in the example cited by Loftin and Ward. If this is the case, we need to know the likelihood of this outcome.

TABLE 9

Estimates of β_1 in the Model with β_1 = 0; w = 121 and 0.1 < p < .9

Estimation Method	Values of ρ	Mean $\hat{\beta}_1$	SDE	MRSE	Range of Estimates Min	Max
MLE	.1	-0.0024	0.0399	0.0372	-0.088	0.0757
	.3	-0.0027	0.0402	0.0375	-0.1094	0.0759
	.5	-0.0034	0.0407	0.0378	-0.1110	0.0760
	.7	-0.0045	0.0418	0.0384	-0.1151	0.0760
	.9	-0.0067	0.0466	0.0399	-0.1325	0.0750
QAD	.1	-0.0005	0.0400	0.0385	-0.1238	0.0826
	.3	0.0049	0.0368	0.0388	-0.1087	0.0861
	.5	0.0108	0.0333	0.0391	-0.0910	0.0876
	.7	0.0170	0.0301	0.0395	-0.0739	0.0854
	.9	0.0231	0.0302	0.0412	-0.0764	0.0904
OLS	.1	-0.0062	0.0424	0.0368	-0.1281	0.0761
	.3	-0.0180	0.0512	0.0379	-0.1706	0.0779
	.5	-0.0368	0.0666	0.0412	-0.2427	0.0879
	.7	-0.0745	0.1002	0.0494	-0.3973	0.1279
	.9	-0.2140	0.2326	0.0781	-1.0095	0.3240

We now consider, in sequence, two further sets of simulation runs that address this issue. The first of these has data generated in exactly the same fashion as for the full model, only the parameter β_1 is set to 0. The second simulation has the same generation procedure, only in this instance β_2 rather than β_1 is set to 0. We will describe briefly[21] the results in terms of bias and the standard deviation of the estimates, comparing them to the results shown in Tables 2 and 4 through 7.

The only table we present for this series of runs is Table 9, which gives the estimation results for the parameter, β_1, set to 0. The mean value of β_1 targets on the generating value of 0 (with a very small downward bias) for MLE. The discrepancies between the mean $\hat{\beta}_1$ and the generating value of 0 are larger than the corresponding bias in the full model. For $\rho \leq 0.3$, QAD is on its target value. The upward bias in the QAD estimates is smaller than the corresponding downward bias for QAD in the full model. As in Table 2, this bias increases with ρ. Using the results of the earlier mathematical argument, we can establish the magnitude of the bias for OLS. From equation 12, the following are generated:

	$\rho = .1$	.3	.5	.7	.9
$\hat{\beta}_1$	–0.0046	–0.0167	–0.0358	–0.0751	–.2245
$\hat{\beta}_2$	–0.3010	–0.3080	–0.3249	–0.3614	–.4568
$\hat{\beta}_0$	10.12	10.58	11.49	13.33	17.99

For β_1, the bottom panel of Table 9 corresponds to these figures. For the parameter set to 0, a downward bias appears, but it is much smaller than the corresponding upward bias in the full model.

The standard deviation of the MLE estimate for this parameter configuration is virtually the same as for the full model, at least for $0.1 \leq \rho \leq 0.7$. A small upward drift occurs in the standard deviation that was not present in the full model. The standard deviations of the QAD estimate are down considerably compared to the corresponding QAD values for the full model. In fact, the QAD estimates shown in Table 9 have lower standard deviations than the MLE method. There is another reversal as well: For values of $\rho > 0.1$, the mean reported standard error for the QAD approach is above the actual standard deviation of the estimate, whereas for the MLE approach the mean reported standard errors are below, and underreport, the actual standard deviation of the estimate. Except for lower values of ρ, the QAD range of the estimates is somewhat lower than the MLE range. Thus, in this instance, QAD outperforms MLE. OLS still performs worse than both QAD and MLE. However, the standard deviation in the OLS estimates is unchanged when compared to those of the full model. The mean reported standard errors are also comparable to those in Table 2, with the exception of the case in which $\rho = 0.9$. The range of the OLS estimates is considerably broader than the range of either QAD or the MLE approach. As there is a consistent bias with OLS and as the mean reported standard error underreports the true standard error, there will be a bias in the direction of seeing parameter estimates for β_1 as significant when the generating value is 0.

With minor numerical variations,[22] the results for β_2 are the same as those reported in Table 4. Differences emerge for β_0. For $0.1 \leq \rho \leq 0.7$ the mean estimates of MLE are close to the results shown in Table 5, but the mean estimate for β_0 with $\rho = 0.9$ is closer (10.06) to the target value. Under MLE, the standard deviation of the estimates tends to be higher than in the full model, but the mean reported standard error is somewhat smaller. The two differences are such that, for this model, the mean reported standard error for MLE underreports the actual standard deviation of the estimates. Whereas the standard deviation of the estimates for MLE was constant throughout the range of the values for ρ, they tend to increase here with the values of ρ and the range of the estimates increases. For QAD a downward bias still exists, but it tends to be smaller for values of $\rho \leq 0.3$. The standard deviation of the estimates for β_0 under QAD tends to decrease with ρ, whereas the mean reported standard error increases (the latter being the same pattern as for the full model). Thus for low ρ the QAD method underreports the standard errors although for higher ρ it overreports them.

In terms of estimating the intercept, the performance of OLS is much better in this model than in the full model. OLS shows less bias (see above). The actual standard deviation of the estimates is unchanged and the mean reported standard error is even less than for the full model. Thus, for OLS, the decrease in the magnitude of the upward bias is offset by an increase in the downward bias of the reported standard error.

Differences also occur for the estimation of the variance of the disturbance term. For MLE there is still the downward bias but the magnitude of the bias decreases as ρ increases. Except for $\rho = 0.9$, the standard deviation of the estimates under MLE are about the same as reported in Table 6; similar results are found for the mean reported standard error. For QAD the downward bias is larger than for the full model, but both the standard deviation of the estimates and the mean reported standard error remain the same as shown in Table 6. For the estimation of ω, OLS is improved when compared to the full model. The bias is smaller, the standard deviation is considerably smaller, especially for higher ρ, and the mean reported standard error is also lower (but still underreports the true standard error by a large amount).

The set of simulations in which β_2 was set to zero is much easier to describe: few changes. For OLS, the mean estimates follow closely the theoretically established results (from equation 12):

	$\rho = .1$	.3	.5	.7	.9
$\hat{\beta}_1$	0.3267	0.4008	0.5261	0.7881	1.756
$\hat{\beta}_2$	–0.0006	–0.0022	–0.0055	–0.0145	0.0512
$\hat{\beta}_0$	11.45	15.70	23.71	43.80	159.0

This apart, the data for β_1 are similar to the figures in Table 2. For β_2, of course, the target value is 0 and all estimators target on it (including OLS). The variability measures are changed little. With MLE and OLS the standard deviations of the estimates of the intercept do not change, but for QAD they rise. The mean reported standard errors rise for all methods. The pattern of over- and underreporting is the same as in Table 5. For ω, there is no real change with $\beta_2 = 0$ compared to Table 6 except when $\rho = 0.9$; similar results are found for ρ.

An overall summary of the runs having one zero slope coefficient is provided in Table 10, where the RMSE is reported for each coefficient for the two models with $\omega = 121$. The left panel deals with the case having $\beta_1 = 0$, and the one on the right deals with $\beta_2 = 0$. We consider the former first. For the estimation of β_1, QAD outperforms MLE, with OLS trailing behind both. For β_2, the RMSE criterion orders the methods as MLE, QAD, and OLS, except for $\rho = 0.9$, when QAD performs better than MLE. For $\rho \leq 0.7$, MLE and QAD are very close. For estimating β_0, QAD outperforms MLE for $\rho \leq 0.5$. Except for $\rho = 0.1$, both MLE and QAD outperform OLS. Estimation of the disturbance term variance shows MLE and QAD as very close (except for $\rho = 0.9$) and both outperform OLS. With the model having $\beta_2 = 0$, MLE is clearly superior to QAD and OLS for estimating β_1; MLE and QAD are close when estimating β_2, but QAD does have a

TABLE 10
Root Mean Squared Errors (RMSE) for Models with one Zero Slope Coefficient: $0.1 \leqslant \rho \leqslant 0.9$ and $\omega = 121$

		Regime with β_1=0			Regime with β_2=0		
Parameter	ρ	MLE	QAD	OLS	MLE	QAD	OLS
β_1	.1	0.0400	0.0400	0.0428	0.0403	0.0821	0.0493
	.3	0.0403	0.0371	0.0543	0.0405	0.0879	0.1117
	.5	0.0408	0.0350	0.0761	0.0406	0.0989	0.2348
	.7	0.0420	0.0346	0.1246	0.0406	0.1043	1.4989
	.9	0.0471	0.0380	0.3161	0.0404	0.0842	1.4850
β_2	.1	0.0568	0.0579	0.0573	0.0569	0.0570	0.0573
	.3	0.0570	0.0577	0.0602	0.0578	0.0566	0.0595
	.5	0.0571	0.0577	0.0702	0.0569	0.0560	0.0655
	.7	0.0576	0.0583	0.0997	0.0569	0.0555	0.0804
	.9	0.0679	0.0580	0.1948	0.0570	0.0558	0.1306
β_0	.1	2.83	3.12	2.91	2.79	4.15	3.32
	.3	2.85	2.93	3.57	2.79	4.53	7.69
	.5	2.89	2.74	4.95	2.79	5.44	14.55
	.7	2.97	2.60	8.07	2.78	6.54	15.57
	.9	3.82	2.62	21.48	2.74	7.56	14.94
ω	.1	20.04	20.46	20.68	20.01	20.86	20.88
	.3	20.02	20.44	24.65	20.02	20.63	25.84
	.5	20.00	20.58	45.16	20.00	20.36	53.16
	.7	20.41	20.78	118.23	20.00	20.09	182.46
	.9	68.99	20.74	527.15	20.00	20.17	171.58

slightly smaller RMSE; MLE is superior for estimating β_0; and both are close for estimating ω, with MLE having the slightly smaller RMSE. In the main, OLS trails, with only two exceptions.

Inference

We now consider inference again. Table 11 shows the counts of correct inference decisions when the null hypothesis is H_0: $\theta = \theta_{gen}$,

TABLE 11
Count of Correct Inferences for $H_o : \theta = \theta_{gen}$

		Model											
		β = (10, 0.3, 0)						β = (10, 0, -0.3)					
		ω = 81			ω = 121			ω = 81			ω = 121		
Parameter	ρ	MLE	QAD	OLS	MLE	QAD	OLS	MLE	QAD	OLS	MLE	QAD	OLS
β_1	.1	99	89	81	98	86	83	96	95	94	96	95	92
	.3	99	91	21	99	89	33	96	97	83	96	98	86
	.5	99	86	0	99	84	2	97	99	70	97	99	74
	.7	99	81	0	99	79	0	97	98	54	97	99	58
	.9	99	87	0	99	86	0	92	97	28	93	98	45
β_2	.1	93	93	94	93	93	94	93	94	94	93	94	94
	.3	93	93	95	94	94	95	93	96	96	93	95	94
	.5	93	93	96	93	93	95	93	95	93	93	95	94
	.7	93	93	98	93	93	98	93	95	84	94	95	89
	.9	93	94	100	93	93	100	92	94	71	91	95	77
β_0	.1	93	95	87	96	94	88	93	93	91	93	93	91
	.3	93	96	33	98	94	41	93	94	87	93	94	87
	.5	93	92	1	100	91	5	93	95	72	93	96	71
	.7	93	90	0	100	87	0	93	96	59	93	97	72
	.9	100	91	0	100	89	0	90	98	43	91	97	40
ω	.1	93	94	94	93	94	94	94	93	94	94	93	94
	.3	93	95	96	93	95	96	94	94	97	94	94	97
	.5	93	95	68	94	95	77	94	96	82	94	95	85
	.7	93	93	3	94	94	8	94	94	21	94	93	29
	.9	93	95	0	94	94	0	91	95	0	92	94	1

(ω = 81, 121 with either β_1 of β_2 set to zero)

where θ_{gen} denotes the generating value of the parameter, θ. The significance level throughout has been taken as 0.01. Of particular interest are those results for $\theta_{gen} = \beta_1 = 0$ and $\theta_{gen} = \beta_2 = 0$. Four parameters are considered in Table 11 (β_1, β_2, β_0, and ω) and two models, each with a zero slope parameter, for which there were runs for $\omega = 81$ and $\omega = 121$ leading to the four vertical panels in Table 11.

Of interest is the extent to which inference decisions based on any of the three methods of estimation are incorrect. Our attention will be given primarily to the slope parameters as decisions concerning them have substantive importance. Consider first the two right-hand panels of Table 11. As the data generated by these models came from a regime where the first slope parameter, β_1, was set to 0, we start by considering the inferential decisions concerning the presence of the corresponding variable in the equation. For $\rho = 0.1$, with the disturbance term variance set at 81, no differences appear between the decisions made by each of the estimation procedures. Roughly five times in a hundred, a mistaken inference will occur under all three methods. However, once the value of the network autocorrelation parameter moves away from this low value, the performance of OLS deteriorates steadily. For $\rho = 0.3$, there are already 17 mistaken inferences in 100, whereas for $\rho = 0.5$ there are 70. For higher values of ρ the performance of OLS is even worse. The performance of MLE is not affected by increases in ρ, except for $\rho = 0.9$, when 8 mistaken inferences occur out of 100. QAD fairs well; if anything, it performs slightly better for higher values of ρ. Much of the same pattern holds when $\omega = 121$: OLS steadily deteriorates, and MLE and QAD perform robustly. The message is clear: Even with small to moderate amounts of network autocorrelation, OLS should not be used to make inferences concerning the presence of X_1 in the linear equation. For configurations such as this, the kind of incorrect inference detected by Loftin and Ward (1983) is very likely. QAD and MLE both perform well, so the insignificance of population density for fertility rates reported by Loftin and Ward (when they use methods designed to deal with network autocorrelation) is consistent with this part of the simulation. Staying within the same regime of data generation, we see that inference concerning the parameter not set[23] to 0 (β_2) is also affected, though to a lesser extent. For moderate ρ ($\rho \leq 0.5$) OLS does perform comparably to both MLE and QAD. However, for higher values of network autocorrelation, the quality, or accuracy, of the OLS inferences begins to deteriorate. On the other hand, the performance of MLE and QAD is robust throughout the range of ρ. However, to the extent that most instances of empirical network autocorrelation do not have values of ρ higher than 0.5, it could be argued that OLS performs well for the lower values of ρ and this is sufficient (for a configuration relatively immune to this kind of network autocorrelation).

The same general pattern is repeated for inferences concerning the intercept: Both MLE and QAD perform robustly throughout the range of ρ, but OLS leads to many incorrect inferences as the network parameter rises. It should be noted that, if anything, QAD performs better than MLE, at least in terms of the count kept of correct inferences. It should also be noted that MLE seems more affected at the extreme value of $\rho = 0.9$. Finally, for the right-hand panel of Table 11, when we consider inferences concerning ω we see the same pattern. The number of incorrect inferences with OLS rises dramatically with increases in ρ while the number of incorrect inferences under MLE and QAD remains fixed.

All in all, for this particular regime of data generation, any inference based on ordinary least squares seems fraught with hazard. This is the case especially when inferences are made concerning the presence or absence of the X variable for which the generating parameter is set to 0 and network autocorrelation is present. Inference concerning the true value of the other slope parameter in this regime is not threatened, at least for small or moderate values of ρ.

Now consider the left-hand panel of Table 11, where β_2 is the parameter set to 0. The results here are more dramatic. Even though β_2 was set to 0, the first horizontal panel of Table 10 makes it clear that inference concerning the value of β_1 has been devasted for OLS. Even for $\rho = 0.1$, OLS performs poorly, leading to approximately 18 incorrect inferences in 100. By contrast, MLE has 1 or 2 incorrect inferences in 100, and QAD has about 12 incorrect inferences in 100. A clear ordering can describe the count of inference decisions for $\rho = 0.1$: MLE is preferable, QAD is suspect,[24] and OLS is very suspect. Poor as the OLS performance is for $\rho = 0.1$, it deteriorates markedly as ρ increases. Already for $\rho = 0.3$, there are 79 incorrect inferences in 100 when $\omega = 81$ and 67 incorrect inferences in 100 when $\omega = 121$. By the time we reach $\rho = 0.5$, in the first horizontal panel of Table 11, we find OLS incapable of returning a correct inference. Although the performance of OLS deteriorates with increases in ρ, both MLE and QAD perform robustly, returning the same number of correct and incorrect inferences as they did for $\rho = 0.1$. The evidence in this panel further suggests that MLE is preferable to QAD.

When we consider inference about the parameter (β_2) set to 0, we find no problems with any method. All three methods perform well and their performance is robust across increasing values of ρ. If anything, OLS performs better than either MLE or QAD, with the suggestion that it performs better with increases in ρ (at least with this parameter in this regime). For both the intercept, β_0, and ω, MLE and QAD perform robustly and adequately. With MLE, mistaken inference results 7 times in 100 and for QAD the result is about the same, although there is less consistency for QAD through different values of ρ. For these parameters, OLS is again inadequate.

Qualitatively, it is easy to summarize these results. Even though OLS does not lead inevitably to the inferential inclusion of a variable in a linear relation when in truth it does not belong there, it is still prone to do so. To the extent that there is interest in the precise parameter values it would appear, from the results shown in Table 11, that OLS is very unreliable. If at some point in the future we are concerned with testing models where the null hypotheses are more interesting than a hypothesis specifying a zero value for a parameter, it will be necessary to have reliable estimation procedures in addition to accurate estimates of the parameters. Both objectives are compromised in a major way by reliance on ordinary least squares. It is more difficult to state a recommendation concerning the relative merits of MLE and QAD. In many instances they do perform comparably, and much of the time QAD is perfectly acceptable. However, given that there are occasions where it does perform less effectively than MLE, a conservative recommendation is to use the maximum likelihood procedure rather than this quick and dirty one.

ESTIMATING THE NETWORK AUTOCORRELATION PARAMETER

When attention is focused on the ability of MLE, QAD, and OLS to estimate β and ω, the simulation results show OLS as clearly inferior. Overall, MLE performs better than QAD, but in many situations they perform comparably. Moreover, there are instances in which QAD performs better than MLE (but in these instances the difference is always small). One variable affecting some of the properties of these estimators is the extent of network autocorrelation. As MLE and QAD both provide estimates of ρ, we consider how well they do this.

We have evidence already in Tables 1 and 7. As shown in Table 1, both MLE and QAD target on $\rho = 0$ when there is no network autocorrelation, with MLE being slightly closer than QAD. In terms of the standard deviation of the estimates, MLE is vastly superior to QAD: Its SDE is 0.027, compared to 0.222 for QAD. Obviously, the RMSE for MLE will be much smaller than for QAD, as both components of the root mean square error are smaller for MLE. Although the standard deviation of the estimate for QAD is 0.222, the mean reported standard error (MRSE) is 0.196. Having an estimator underreport the extent to which it varies across all possible estimates—given a particular parameter configuration—is undesirable. By contrast, having MLE overreport the MRSE can be viewed as an improvement over QAD. Unfortunately, the extent to which MLE overreports its sample variability is considerable, with the RMSE almost 6 times as high as the SDE. Although avoiding cavalier inferences, this surely is laying the foundation for an inference that is much too conservative.

When regimes involving network autocorrelation are considered (see Table 7), the same pattern is evident.[25] The standard deviation of the estimates, for each value of ρ, is much higher for QAD than MLE. Although the mean reported standard error of QAD is smaller than the SDE for small value in ρ and larger for higher values of ρ, it is always relatively close to the actual standard deviation of the estimates. This is not so for MLE, where there is, for all ρ, a wide gulf between the true SDE and the established MRSE.

The foregoing estimation properties of MLE (i.e., bias and the standard deviation of the estimates) indicate that it is superior to the QAD approach. Except for $\rho = 0.1$, the MLE estimate is always closer to the generating value of ρ used in the full models. Consistent with the smaller standard deviation of the estimate, the range of values for the MLE estimate is much smaller. The MLE estimates target tightly around their central value, and the QAD estimates spread widely. For both estimators, the variability across runs diminishes with increasing values of ρ. For the smaller values of ρ the quick and dirty approach provides alarmingly wide ranges for its estimates of the parameter. It is only for values of ρ in excess of 0.5 that the variability of the QAD estimates can be viewed as reasonable. From the evidence contained in Table 7 it is clear also that MLE is preferable to QAD for estimating the value of the network autocorrelation parameter. However, we can push this further and look more closely at the consequences for the actual inference decisions made when using these estimation strategies.

Table 12 provides further evidence for the superiority of MLE. There are two panels in Table 12; the first deals with inference for the full models and the second with inferences when one of the slope parameters has been set to 0. This table provides the count of correct inferences for the null hypothesis specifying no network autocorrelation. Qualitatively, we need to distinguish the situations in which there is no network autocorrelation from those in which there is such an effect. The first row of the upper panel gives the correct nonrejections of this null hypothesis when no network autocorrelation was used to generate the data. Without exception, the MLE procedure leads to correct inferential decisions, although this is not true for QAD, which rejects this null hypothesis between 5 and 8 times in 100. There may be several reasons for this difference. First, some of the incorrect decisions under QAD may occur simply by chance. After all, the number of incorrect rejections of the null hypothesis is around some of the conventional significant levels used. A second factor may be the great variability in the QAD estimates. As there are extreme

TABLE 12
Counts of Correct Inference for $H_0 : \rho = 0$ for Regimes with Differing Disturbance Term Variance, ω

Disturbance Term Variance and Estimation Method

Values of ρ	ω = 25 MLE	ω = 25 QAD	ω = 49 MLE	ω = 49 QAD	ω = 81 MLE	ω = 81 QAD	ω = 121 MLE	ω = 121 QAD
0	100*	95*	100*	96*	100*	94*	100*	92*
0.1	0	17	0	11	0	12	0	6
0.3	100	93	100	78	100	70	98	66
0.5	100	100	100	92	100	100	100	99
0.7	100	100	100	90	100	100	100	100
0.9	100	100	100	100	100	100	100	100

(a) Full Models

Model, Disturbance Term Variance and Estimation Method

Value of ρ	$\beta_2=0$ ω = 81 MLE	$\beta_2=0$ ω = 81 QAD	$\beta_2=0$ ω = 121 MLE	$\beta_2=0$ ω = 121 QAD	$\beta_1=0$ ω = 81 MLE	$\beta_1=0$ ω = 81 QAD	$\beta_1=0$ ω = 121 MLE	$\beta_1=0$ ω = 121 QAD
0.1	0	13	0	13	1	14	0	13
0.3	100	61	100	42	100	58	100	58
0.5	100	100	100	98	100	91	100	67
0.7+	100	100	100	100	100	100	100	76

(b) Models with One Slope Parameter Zero

+ For $\rho = 0.9$ for $\beta_i = 0$, each cell entry is 100

*For $\rho = 0$ the correct decision is to not reject H_0; otherwise the correct decision is to reject H_0.

estimates (see Table 1), it would be surprising if they did not lead to the rejection of this null hypothesis. Third, the large over-estimate under MLE of the sample variability means that even the extreme estimates are not large enough—in relation to the very conservative estimate of the standard error—to lead to rejection of this null hypothesis.

Panel a of Table 12 suggests that not until $\rho = 0.3$ is it possible for either estimation strategy to detect values of ρ significantly different from 0. For $\rho = 0.1$ neither method leads to frequent rejection of the incorrect null hypothesis: MLE never does so, and QAD fares a little better. For such a small value of ρ this should not be too surprising. For $\rho = 0.3$, a value large enough to permit frequent rejections of the wrong null hypothesis $\rho = 0$, we see a considerable difference between MLE and QAD. For $\rho \leq 0.3$ there is only one instance in which the MLE does not return 100 correct inferential decisions out of 100: for $\omega = 121$, and $\rho = 0.3$. Otherwise, the MLE estimation method is always correct with regard to this particular inference.[26] The QAD performance is always worse relative to MLE, and QAD decisions deteriorate with increases in disturbance term variance.

A similar pattern appears when we consider the models generating data having one of the slope parameters set to 0. These results are shown in panel b of Table 12. For $\rho = 0.1$, the MLE approach does not permit discriminating this value from 0, whereas the QAD does around 13 times in 100. Once ρ reaches this threshold value of 0.3, the MLE approach tends to support the inference that the true value of ρ is not 0. By contrast, the QAD approach permits only the correct inference between 42 times in 100 and 61 times in 100. Even by the time the network autocorrelation parameter reaches 0.5, there are still instances with β_2 set to 0 in which the QAD approach does not permit the rejection of the incorrect null hypothesis. Undoubtedly, this is due to some of the extremely low estimates returned by this method. When β_1 is set to 0 the performance of QAD is completely inadequate in correctly rejecting the null hypothesis specifying no network autocorrelation. Even for ρ as high as 0.7, with disturbance term variance of 121, there are still 24 incorrect decisions in 100 under QAD.

In summary, QAD is largely inadequate for determining the presence of network autocorrelation. This was suggested in panel a of Table 12 for moderate values of ρ and is emphasized in panel b of Table 12 more or less for all values of ρ. Obviously, it is difficult to generalize from a restricted number of simulation runs, but the evidence here points to the inadequacy of QAD. It should be pointed out, again, that MLE looks perhaps a little better than it should because of the tendency to overestimate the sample variability of $\hat{\rho}$.

TABLE 13
Count of Correct Inference for H_0: $\rho = \rho$ gen for QAD: $\omega = 81$, $\omega = 21$

	Disturbance Term Variance ω			
Values of ρ	25	49	81	121
0.1	95	89	94	87
0.3	97	94	93	91
0.5	94	92	91	88
0.7	92	90	86	80
0.9	91	88	87	86

(a) Full Models

	Model and Disturbance Term Variance			
	$\beta_2 = 0$		$\beta_1 = 0$	
Values of ρ	$\omega = 81$	$\omega = 121$	$\omega = 81$	$\omega = 121$
0.1	93	92	87	92
0.3	88	88	88	88
0.5	80	80	83	80
0.7	82	76	73	76
0.9	87	86	82	86

(b) Models with One Zero Slope Parameter

Because of this tendency to overestimate the variability of $\hat{\rho}$, the MLE approach always leads to the correct inference when the null hypothesis specifies $\rho = \rho_{\text{gen}}$. By contrast, Table 13 displays the count of the correct inferences for this null hypothesis for QAD method. Panel a of Table 13 shows the counts for the full models, and panel b shows the counts where one of the slope parameters has been set to 0. The full models show that as the disturbance term variance increases, the number of correct inferential decisions with QAD decreases. The models with one zero slope parameter show that as the network autocorrelation increases, the tendency is for the number of correct decisions to decrease. In summary, with the exception of $\omega = 25$, the QAD approach is not accurate in correctly accepting the null hypothesis specifying the generating value of ρ.

In estimating ρ and determining whether or not there is a network autocorrelation problem to be addressed, QAD is clearly inferior to MLE. This, together with the Monte Carlo results presented above, suggests that the quick and dirty approach does not suffice when one is considering models with network autocorrelation. Although MLE is superior to QAD, there is the nagging problem that, as an estimation method, it badly overestimates the sample variability of ρ. Ideally, it would be useful if a legitimate method of deflating this tendency could be established.

CONCLUSION

As MLE is the preferable estimation strategy, our recommendation is to use it on all occasions rather than QAD or OLS. We are less certain about a recommendation concerning the overestimation of the variability of $\hat{\rho}$ under MLE. If the MLE estimate of ρ is equal to or above 0.3, then we feel confident that the inference concerning the presence of network autocorrelation as a "network effect" will be correct. For $\hat{\rho} \leq 0.1$, the inference will state (correctly, in the main) that no network autocorrelation effect is present. For $0.1 \leq \rho \leq 0.3$ the results will be equivocal. At the high end of this range, there will be cases where there is a network effect, but MLE will lead to an inference that there is no such effect. We suggest that MLE be used anyway, so that the estimation of β and ω, and inference concerning their values, will have a sounder basis. Unless a theory predicts a network effect, little will be lost. But when such a prediction is made,[27] the inability of MLE to detect a network effect—when it is present—becomes serious.[28]

We have barely scratched the surface with these simulations. Apart from a broader range of values for the β_1 and different distributions of X, other avenues of exploration appear important. One is to allow variation of N and another is to explore W matrices with different structures. Both have been tackled by Dow et al. (1982) for the network disturbances models, and similar explorations will be important for network effects models.

NOTES

1. More precisely, for $y_i = X_i\beta + \epsilon_i$, the ϵ_i and ϵ_j are assumed to be drawn independently from normal distributions of mean 0 and variance ω: $\epsilon \sim \mathrm{IN}(0, \omega I)$.

2. This example is particularly vivid, as the hypothesis relating fertility and population density is widely accepted, largely because of an early OLS analysis. The evidence in Doreian (1981) of different inferences stemming from either taking or not taking spatial autocorrelation into account is much less vivid.

3. "Network autocorrelation" will be used here as a general term to denote all forms of autocorrelation, other than time-series models, where there are interdependent data points. This includes spatial autocorrelation as a special instance.

4. Doreian (1982) considers the linear model with both a network disturbance term and a network effects model, but such a model is not considered here.

5. Strictly, $\omega = (1/N)\, y'A'MAy$ from the log-likelihood function. However, this is a biased estimator, although equation 6 corrects for this bias.

6. In principle, there are many ways to construct W. We have stayed close to one particular empirical instance of interdependent data points. So, through the runs, N = 64.

7. In Doreian (1981) there are estimators of ω for different data sets of 22.9 (Table 4), 49.78 (Table 2), and 126.4 (Table 1). These are close to 5^2, 7^2, and 11^2, and we added the value of 9^2 to complete the sequence.

8. The total number of tables that could be reported for these simulations is around 140. Obviously, they cannot all be included. So we present some and outline the form of the remaining tables in relation to those included here.

9. It is about 16% above the actual standard deviation (whereas the QAD report is about 21% below its true value, and OLS is 7% below its true value).

10. We would like to claim this was built into the design but, in fact, it was fortuitous. It is reasonable that some configurations are impervious to network autocorrelation and others are sensitive. Both are featured into one (generic) set of exogenous variables used in this simulation.

11. For example,

	ω	25	49	81	121
(MLE)	$\sigma(\hat{\beta}_1)$	0.018	0.026	0.033	0.040
	$\hat{\sigma}(\hat{\beta}_1)$	0.026	0.034	0.041	0.047

12. There is a slight upward drift in the bias from 0.0002 (for $\rho = 0.1$) to 0.0044 for ($\rho = 0.9$), but it seems inconsequential relative to the magnitude of the coefficient and its estimates.

13. Again, there is a slight upward drift in the mean reported standard error from 0.0482 to 0.0551. This too seems inconsequential, although in other empirical situations this may not be the case.

14. Although the decrease from 0.0690 to 0.0638 is probably inconsequential also, it is interesting that the variability in $\hat{\beta}_1$ decreases with ρ.

15. There is a slight upward drift in mean ($\hat{\beta}_2$) with MLE but, like earlier drifts, it is numerically small for this model with this regime.

16. If the regression is restricted such that the fitted plane passes through the origin, we have MRSE = 1.23 SDE + 0.0088 ρ. Note that each pair of data values has a corresponding value of ρ. For this analysis, ρ is a variable, ranging across runs, and is not a constant. Also, as each run is completely independent of every other run and batches of 100 runs are independent of other batches, we do not need to take into account network autocorrelation in these crude summaries.

17. The intercept of –0.00004 is ignored.

18. With the data ordered by SDE, a regression of MRSE on SDE has a low Durbin-Watson statistic suggesting curvilinearity.

19. We detail consideration of ρ later in the article.

20. One exception to the comparable quality of MLE and QAD can be discerned for the parameter β_1, where, perhaps, MLE is marginally superior to QAD. One exception to the inferiority of the OLS performance is found with the estimates concerning β_2. But, in general, the message is clear: OLS is very unreliable and, thus far, MLE and QAD are comparable in inference decisions concerning the actual parameter generation values.

21. With tables corresponding to Tables 2 and 4 through 7 for each of these runs, we have too many to present here and resort to a verbal summary, together with a table of root mean square errors.

22. Of course, the bias for OLS follows the pattern described earlier in this section.

23. This is with H_0: $\theta = \beta_2 = -0.3$, the kind of hypothesis that is unlikely to be tested, as theories are seldom powerful enough to instate a nonzero parameter value. Even so, if such a hypothesis were offered, it would be rejected 7 times in 100 even when true.

24. For QAD, mistaken inference (at $\alpha = 0.01$) ranges between 9 and 14 times in 100. The declared value of $\alpha = 0.01$ may not actually be the one operative.

25. The RMSE for the estimates of ρ, corresponding to the data given in Table 7, are as follows:

	.1	.3	.5	17	19
MLE	0.0263	0.0283	0.0173	0.0112	0.0051
QAD	0.2087	0.1858	0.1661	0.1257	0.0497

26. This is not solely a virtue, as this is due partly to an overestimate of the sample variability.

27. Such instances do exist (see Burt and Doreian, 1982).

28. Our attempts to generate a relation between SDE and MRSE for the MLE method have not been successful. Rather, they do not lead to useful advice. Combining the runs for all the full models for $\rho > 0$, we can establish via regression methods that

$$\underset{}{\text{MRSE}} = \underset{(0.005)}{0.053} + \underset{(0.438)}{6.38}\ \text{SDE} - \underset{(13.17)}{111.0}\ (\text{SDE})^2 - \underset{(0.005)}{0.040}\ \rho \qquad R^2 = .99$$

Then for the models with a slope parameter set at zero, a similar operation yields:

$$\underset{(0.008)}{\mathrm{MRSE} = 0.134} + \underset{(0.693)}{2.445}\ \mathrm{SDE} - \underset{(12.63)}{38.28}\ (\mathrm{SDE})^2 - \underset{(0.008)}{0.124}\ \rho \qquad R^2 = .97$$

Such instability across empirical conditions cannot provide the basis for a policy of adjusting the MLE standard errors. In part, this may be due to the collinearity of SDE and $(\mathrm{SDE})^2$. Although ridge regression may reduce some of this instability, we doubt that this is the best avenue for pursuing the relation between SDE and MRSE for the maximum likelihood method.

REFERENCES

BURT, R. S. and P. DOREIAN (1982) "Testing a structural model of perception: conformity and deviance with respect to journal norms in elite sociological methodology." Quality and Quantity 16: pp. 109-150.

DOREIAN, P. "Maximum likelihood methods for linear models." Soc. Methods and Research 10 (February): 243-269.

———(1981) "On the estimation of linear models with spatially distributed data," pp. 359-388 in S. Leinhardt (ed.) Sociological Methodology. San Francisco: Jossey-Bass.

———(1980) "Linear models with spatially distributed data: spatial disturbances or spatial effects?" Soc. Methods and Research 9 (August): 29-60.

DOW, M. M., M. L. BURTON, and D. R. WHITE (1982) "Network autocorrelation: a simulation study of a foundational problem in regression and survey research." Social Networks 4: 169-200.

GOLDBERGER, A. S. (1964) Econometric Theory. New York: John Wiley.

LOFTIN, C. (1972) "Galton's problem as spatial autocorrelation: comments on Ember's empirical test," Ethology 11: 425-435.

———and S. WARD (1983) "A spatial autocorrelation model of the effects of population density on fertility." Amer. Soc. Rev. 48 (February): 121-128.

ORD, K. (1975) "Estimation methods for models of spatial interaction." J. of the Amer. Stat. Assn. 70 (March): 120-126.

RAO, P. and R. L. MILLER (1971) Applied Econometrics. Belmont, CA: Wadsworth.

WHITE, D. R., M. L. BURTON, and M. M. DOW (1981) "Sexual division of labor in African agriculture: a network autocorrelation analysis." Amer. Anthropologist 83 (December): 824-849.

Patrick Doreian is Professor of Sociology and has interests in network analysis, methods for spatial or network autocorrelation, and network approaches to dependency theory and societal evolution.

Klaus Teuter is a Programmer at the Social Science Computer Research Institute. He has an M.A. in sociology. Currently he is preparing SSCRI's participation in Pitt's computerized Campus of the Future.

Chi-Hsein Wang is a Ph.D. student at the Graduate School of Public and International Affairs and Programmer at the Social Science Computer Research Institute at the University of Pittsburgh.