THEORY, MEASUREMENT, AND SPECIFICATION ISSUES IN MODELS OF NETWORK EFFECTS ON LEARNING*

Reply to Carbonaro and to Hallinan and Kubitschek

Stephen L. Morgan
Harvard University

Aage B. Sørensen
Harvard University

Our article (Morgan and Sørensen 1999, henceforward M&S) is an attempt to spark debate and encourage research on the effects of school social systems on learning. The comments by Carbonaro (1999, henceforward C) and Hallinan and Kubitschek (1999, henceforward H&K) suggest that we have succeeded. We are thankful for C's appraisal, supplementary data analysis, and comparison to his own related research (Carbonaro 1998). We are, however, surprised by H&K's comment. As we will show, H&K do not understand Coleman's paired theories of norms and social capital, which we embrace and extend in our article. Thus, much of H&K's conceptual criticism, although directed at us, applies to Coleman's theories, which H&K's comment aims to defend. In response to the comments of C and of H&K, we offer a defense of our article, starting with two general comments.

First, our distinction between norm-enforcing and horizon-expanding schools is an ideal-type distinction, constructed primarily to ease the presentation of our findings.[1] The distinction could clearly be differentiated further, and at first we considered a fourfold schema, selecting local high schools from the greater Boston area as our examples. Thus, we find C's three-fold typology of nonfunctional, dysfunctional, and functional school communities useful. We regard it as theoretically consonant and therefore supportive of our norm-enforcing and horizon-expanding distinction and the more fine-grained claims we make throughout our article.

Second, we all agree that the NELS data are imperfect, but they are the best we have and are worthy of careful analysis. The questions on parental networks incorporated in the NELS survey were developed explicitly to evaluate Coleman's hypotheses about the possible beneficial effects of social closure on student outcomes. H&K claim that the egocentric network data collected from parents should properly be considered as measures of average student sociability rather than as the properties of parental networks. H&K seem unaware that the same measures can serve as operationalizations of different concepts. Furthermore, they do not attempt to show that their alternative conceptualization of the measures can account for the findings we present, but merely assert that our conceptualization is wrong—and therefore that our findings are irrelevant.

H&K's criticism of our decision to analyze these network data, in an attempt to evaluate Coleman's hypotheses, using measures that he helped develop, implicitly encourages scholastic retrogression that cannot be good for sociology. Their critique accuses us of

* Direct all correspondence to Stephen L. Morgan, Department of Sociology, Harvard University, William James Hall 5th Floor, Cambridge MA 02138 (smorgan@wjh.harvard.edu), or Aage B. Sørensen at Department of Sociology, Harvard University, William James Hall, Cambridge MA 02138 (abs@wjh.harvard.edu).

This research was supported by a fellowship for Morgan from the Spencer Foundation and by research grants to Sørensen and to Morgan from the American Educational Research Association, which receives funds for its AERA Grants Program from the National Science Foundation and the National Center for Education Statistics (U.S. Department of Education) under NSF Grant #RED-9452861. Opinions reflect those of the authors and do not necessarily reflect those of the granting agencies.

[1] Our article advances a very modest theory of network effects on learning, primarily because it is an empirical article constrained by available data. A more complete version of this theory is presented in Morgan (1998).

three apparently unforgivable sins: We used the data for the purposes for which they were collected; we placed faith in the instrument that the survey designers crafted; and we took seriously Coleman's social capital explanation for school effects.

THEORETICAL ISSUES

H&K question our distinction between norm-enforcing schools and horizon-expanding schools. First, they argue that we do not recognize the potential harm that the social structures surrounding our ideal-type norm-enforcing school can cause students. They make this claim even though the possibility of such harm is the main point of our article.[2] Perhaps worse, H&K seem unaware of the details of Coleman's social capital theory of school effects. H&K argue that "to call a school norm-enforcing with respect to academics because it contains a high degree of intergenerational social closure seems inappropriate. . . . Friendships have different bases of attraction, many of which are unrelated to education" (p. 687). This statement suggests that H&K are not only unfamiliar with our article but also with Coleman's theoretical argument about the dependence of the creation of norms on preexisting social networks. Therefore, we draw much of our defense against H&K's criticisms from Coleman himself.

Coleman's major project for the last 30 years of his life was the construction of macro-social concepts on the micro-foundations of rational choice theory. One of the most important accomplishments of this project was his development of a theory of norms (see Coleman 1990, chaps. 10–11). As is clear from the theory, networks are essential for the realization of effective norms because closed loops in networks create opportunities for the imposition of sanctions on those who violate norms. It does not matter why or how these networks are formed. Ties can be based in interpersonal attraction, work interaction, bowling league membership, or whatever. Coleman (1990) writes of the ". . . second condition for emergence of an effective norm . . . [where] beneficiaries of a norm, acting rationally, either will be able to share appropriately the costs of sanctioning the target actors or will be able to generate second-order sanctions" and then claims that this condition crucially "depends on the existence of social relationships among beneficiaries" (p. 273). Coleman goes on to identify the network properties that support the creation of effective norms, and he then proceeds to develop the concept of social capital.

For Coleman, effective norms are an essential ingredient of the social capital available to schools. The abundance of such social capital in some schools, particularly in Catholic schools, creates more learning than in less endowed schools because student behavior is regulated by strict achievement norms. Coleman's argument does not in any conceivable manner depend on whether the networks sustaining these norms are education relevant.

Coleman's theory of norms is not the only piece of his writing that is apparently unfamiliar to H&K. They further argue that it is "not logically sound" (p. 688) to assume a negative correlation for parents, either as individuals or in the aggregate, between ties to adults within a community and ties to adults outside of a community. This assumption, however, is Coleman's own assumption, not ours, and it is probably why he encouraged the NELS investigators to collect the network data in the way they did. In describing one of his sociograms in his chapter on social capital, Coleman (1990) writes: "If A and B are adults in a community . . . , then closure in the community can be pictured as . . . arrows from one actor to another. . . Lack of closure is shown . . . [by the absence of arrows in the second panel of the sociogram] where the parents, A and B, have their friends outside this community" (p. 318). In evaluating and extending Coleman's hypotheses, we maintain his implicit assumption.

H&K and C both argue that social closure should be regarded as an individual-level at-

[2] For example, H&K write: "Parents and students in socially closed networks may have norms that are unrelated to academics, favorable to academics, or hostile to school norms and practices. Thus, a school with a high incidence of intergenerational social closure is not necessarily characterized by shared parental norms about academic interests and concerns" (p. 687). We fail to see how this can be regarded as an objection to the theoretical argument of our article.

tribute rather than as a school-level attribute. We view social closure as a feature of social structure that is rightly regarded as a group-level attribute that yields returns to individuals.[3] Again, we rely on Coleman for our defense. In a section titled "The Public-Good Aspect of Social Capital," Coleman (1990) writes: "As an attribute of the social structure in which a person is embedded, social capital is not the private property of any of the persons who benefit from it" (p. 315).[4] Nonetheless, for comparison with C's supplemental models, we report models in Table 1 (on p. 698) that estimate both school-level and individual-level effects of *friends in school* and *parents know parents* on mathematics achievement gains.

We outline in our article the possible mechanisms by which social capital in its various forms can enhance learning and the mechanisms by which it can inhibit learning. In this we go further than Coleman, who only hints at the dark side of social capital. None of the conceptual criticisms leveled by H&K has any bearing on these mechanisms or on the empirical results they produce.

Not only do H&K dislike our extension of Coleman's theory, they inadvertently demonstrate their own dislike of Coleman's theory. The "conceptual ambiguity" that they attribute to us rests on their misunderstanding of the mechanisms and assumptions that Coleman maintains in his writing and that we adopt. Clearly, we like Coleman's theory and our extension of it much better than do H&K, and we are grateful for the opportunity to defend them both.

CRITICISM OF METHODOLOGY

Some Variables Measure the Same Concepts

H&K level at us a serious methodological charge, claiming that two variables—*friends in school* and *parents know parents*—measure the same concept—the average sociability of students in a school. Coleman urged that the NELS parent questionnaires measure the network relations among parents. If the questions used to construct *friends in school* and *parents know parents* were intended to be measures of student sociability rather than parental connectedness, they would not have been included on the parent questionnaire.[5] H&K provide "evidence" for their criticism in a simulation of our variable construction that yields correlations among variables that depart substantially from both our own and C's estimates. As claimed by H&K, the average number of friends named by a parent in a school is indeed positively correlated with both *friends in school* and *parents know parents*—.598 and .512, respectively, by our estimates. These correlations have little or no relevance for our analysis unless one assumes, as H&K apparently do, that the NELS network measures can only be used to measure sociability and not the social capital that is generated by social networks as hypothesized by Coleman.

Multicollinearity and Parameter Specification

H&K then argue that models that specify separate effect parameters for both *friends in*

[3] C claims that our specification of social closure as a school-level attribute requires that "students with low levels of individual closure who attend a school with a high *overall* level of closure will benefit as much as students with high individual levels of closure" (p. 682, italics in original). We disagree. Just as with any other public good, returns harvested from social closure and any norms or information that it supports can vary over individuals in a given community. Our models indicate that the school *means* of such individual-level returns vary meaningfully with variations in social structures.

[4] Directly related, Coleman provides a perfect example for us that is also a defense against H&K's objection that closure must exist among the vast majority of parents in a community for social capital benefits to accrue. Coleman (1990) writes: "For example, where there exists a dense set of associations among some parents of children attending a given school, these involve a small number of persons, ordinarily mothers who do not hold full-time jobs outside the home. Yet these mothers themselves experience only a subset of the benefits of this social capital generated for the school" (p. 316).

[5] The questionnaire asks parents: "Please list the first names (or nicknames) of your teenager's close friends and indicate: (A) whether the friend attends school with your teenager, and (B) whether you know the parent/s of that teenager."

school and *parents know parents* are nonsensical because their multicollinearity produces misleading partial regression coefficients. With our variable construction, the correlation between *friends in school* and *parents know parents* is either .484 or .510, depending on whether one estimates a weighted correlation of the school means for the 898 schools or the weighted correlation of the school means for the 9,241 students in those schools.[6]

Multiple regression analyses would rarely be useful if there were no multicollinearity among independent variables. A correlation coefficient of .5 for two independent variables is not too large. Clearly, multicollinearity reduces the amount of information available for inference and can therefore affect one's confidence in the apparent meaning of a point estimate. Fortunately, the standard error measures the seriousness of this problem for each coefficient. As is clear from the standard errors of our results, H&K have little or no basis for questioning our interpretation of our models. The basic question worth asking is how sensitive are the point estimates to alternative specifications of the model (Winship 1998)? As we demonstrated in our Appendix B (see p. 678) and in the supplement (available from us on request), we have worried a lot about such issues and regard our point estimates as robust to many alternative specifications.

[6] In contrast, H&K claim that the correlation is .66, analyzing a sample of 10,602 students in 972 schools in order to "illustrate variable construction" (p. 689) rather than replicate our analyses. It is unclear how they obtain such a high correlation, or why they did not select the same sample that we did. The difference in the estimated correlations probably results from their inclusion of students who were used to "freshen" the sample in 1990 and students in other types of private schools. Moreover, C reports a correlation between the two variables of .583—midway between our correlation and H&K's. We presume that C's correlation is estimated for the same sample of students used for his regression that includes twelfth-grade dropouts and students who have transferred between schools—students that we exclude from our analysis. Thus, we regard C's correlation as an overestimate, and H&K's is even larger. Finally, if we discard students in schools with fewer than 10 students, as we did in the regression in column 3 of Table 1 (on next page), the correlation between *friends in school* and *parents know parents* is nearly the same at .505.

Are School-Level Variables Meaningful?

Aside from arguing that social closure should only be considered an individual-level attribute, C further argues that school-level measures of social closure based on NELS data are inadequate because of the small and nonrandom within-school samples that make up the NELS data set. Carbonaro (1998) estimates only individual-level effects of parental ties, which he labels social closure.[7]

In his comment, C estimates models with school-level network variables and individual-level departures from these school-level means. These models serve as a sensitivity check on our analysis. Even if one makes alternative choices about how to handle missing data, decides not to model sample attrition, includes twelfth-grade dropouts in models that estimate math achievement in high school, and adjusts for prior achievement in mathematics in high school with eighth-grade test scores, the same basic parameter estimates for *parents know parents* (which C labels *social closure*) and *friends in school* emerge. These models should not be regarded as extensions of our models but rather as extensions of the models he estimated in Carbonaro (1998).

Table 1 presents two models in response to C's comment. Column 1 presents an ex-

[7] Carbonaro (1998) examines a wide range of outcomes in order to evaluate the reach of the concept of social capital. Carbonaro estimates individual-level effects of parental closure and then adds other independent variables to the models in an attempt to explain them away. Generally, he is unable to explain away the parental-closure effects and therefore concludes that Coleman's hypothesis has some support in the NELS data. However, his analysis is open to challenge. He dilutes the explanatory power of covariates that compete with parental closure because he imputes unconditional means for missing values. He also does not explicitly model sample attrition, does not account for sector effects on learning, and relies on an outdated measure of socioeconomic status. And as we mentioned in footnote 7 in our article (p. 670), his claims about the positive effect of parental closure on math achievement in high school are untenable.

Table 1. Coefficients from the Regression of Math-Score Gains between the Tenth and Twelfth Grades on Selected Independent Variables: Revised Models in Response to Carbonaro (1999) and to Hallinan and Kubitschek (1999)

Independent Variable	In Response to Carbonaro: Model 4 with Individual-Level Effects	In Response to Carbonaro: Model 4 with Individual-Level Effects for School Samples ≥ 10	In Response to Hallinan and Kubitschek: Model 3 with Modified *Social Closure* Variable	In Response to Hallinan and Kubitschek: Model 4 with Modified *Parents Know Parents* Variable
FIXED EFFECTS				
Constant	4.307	4.302	4.317	4.307
School-Level Variables				
Catholic school	1.731***	1.821***	1.499***	1.686***
	(.321)	(.430)	(.326)	(.321)
Social closure around school	—	—	.068	—
			(.113)	
Parents work together	—	—	.194	—
			(.366)	
Parents have adequate say	.537*	.633	.412	.538*
	(.238)	(.357)	(.328)	(.238)
Friends in school	.380**	.497**	—	.554**
	(.125)	(.186)		(.171)
Parents know parents	–.314*	–.601**	—	–.388*
	(.133)	(.206)		(.167)
Student-Level Variables				
Friends in school	.056	.045	—	—
	(.062)	(.071)		
Parents know parents	–.046	–.051	—	—
	(.056)	(.065)		
IRT math score in 10th grade	–.107***	–.109***	–.106***	–.107***
	(.006)	(.007)	(.006)	(.006)
RANDOM EFFECTS				
School-level variance	.890	.975	.929	.896
(.232)	(.265)	(.229)	(.229)	
Student-level variance	28.108	28.609	28.109	28.108
(.852)	(.992)	(.851)	(.852)	
–2 log-likelihood	57,617	44,675	57,627	57,618
Number of schools	898	500	898	898
Number of students	9,241	7,146	9,241	9,241

Notes: Robust standard errors are in parentheses. Data are weighted at both the student and school levels. Additional school-level and student-level covariates are the same as for Table 3 in Morgan and Sørensen (1999). Because the model in column 2 is estimated for a nonrandom subsample of students, and because it is estimated using the same weights as the other models, one cannot infer population distributions from the model in column 3 without additional assumptions that we do not wish to invoke. The *social closure* and *parents know parents* variables were modified for the models in columns 3 and 4 by ignoring the reported ties of students' parents to parents of friends who do not attend the same schools as the students.

*$p < .05$ **$p < .01$ ***$p < .001$ (two-tailed tests)

tension of Model 4 from Table 3 in M&S (p. 669) that includes individual-level variables for *friends in school* and *parents know parents*, as C suggests. The estimated individual-level effects are close to zero and are smaller than their standard errors.

C's main methodological complaint is that the within-school samples of the NELS data are incapable of producing meaningful estimates of school-level network properties. In particular, he argues that the samples are too small and are not simple random samples of the within-school populations. Over the 9,241 students in 898 schools in our analysis sample, the average student is in a within-school sample of 13.53 students and the average school has a sample of 10.63 students. We would, of course, be happier with the NELS data if these within-school samples were larger. However, the question is not whether each school sample is large enough to provide a meaningful estimate of the school-level network properties of each school but rather whether over the whole sample of schools the school-level network properties are on average well estimated. Determining whether this is the case is, by construction, beyond the capabilities of the NELS data. But the small size of the within-school samples cannot be regarded as an a priori justification for discarding school-level measures. Moreover, our models are not of a complexity that requires large within-school samples—we do not model school-level variation in within-school relationships.

The fact that within-school samples are not simple random samples is not a severe problem either. By appropriately weighting for differential sampling probabilities within schools—to adjust for the oversampling of Asian and Hispanic students—when computing school-level means and by explicitly modeling sample attrition by using a control function of propensity scores for inclusion in the analysis sample, we hope to have mitigated potential problems with the within-school generalizability of the within-school samples.

To examine whether our findings are produced by the preponderance of small schools in our samples—and to test C's claim that there may be important individual-level closure effects that are not picked up in small school samples where there are few individual-level departures from the school means of closure—column 2 of Table 1 presents the same model as shown in column 1 but for the 500 of the 898 schools that have 10 or more students.

The coefficients from this model are even more supportive of our interpretations than those presented in our article, suggesting that the smallest within-school samples in the full sample may introduce measurement error that attenuates estimated effects of the network variables. These 500 schools, however, do not represent a sample that can be generalized to any population of schools or students. We do not, therefore, want to claim that the estimates reported in column 2 are better than those reported in our article. Nonetheless, our main point stands: The coefficients are not produced by odd variation or overdispersion in the estimated means of small schools because, as in all Bayesian multilevel models, the coefficients are precision weighted by within-school sample size.

Modified Measures of Social Closure

H&K make one interesting point. They argue that it might be truer to Coleman's original hypothesis to construct a social closure variable that ignores the reported ties of a student's parents to the parents of friends who do *not* attend the same school as the student. Accordingly, we modified the *parents know parents* and *social closure around school* variables, reducing their weighted means from 3.121 and 3.067 to 2.534 and 2.794, respectively. We then reestimated Models 3 and 4 from our Table 3 using these modified variables (see columns 3 and 4 of Table 1). Again, our interpretations stand.

CONCLUSION

Our paper may suffer from weaknesses, but few of those claimed by C and H&K can be counted among them—other than that we may have placed too much faith in the NELS data and its survey designers. Perhaps IRT scaling of the math tests does not solve all the problems it is supposed to. Perhaps our propensity-score modeling of sample selection is not careful enough. Perhaps the technology of multilevel Bayesian models is not developed enough to handle complex survey data. Perhaps linear constraints on the network effects are unreasonable. And perhaps network properties are endogenous and therefore our estimates are contaminated by selection biases. For any of these reasons, our findings, like those of any empirical study, may not be confirmed in a future replication.

Nonetheless, the main conclusion of our article stands without dispute, even among the commentators: With a careful analysis of the best available data, the well-studied Catholic-school effect on mathematics achievement cannot be explained away by any specification of network closure variables. Our secondary conclusion is more controversial, but is consistent with the NELS data and is based on a reasonable theoretical argument. *Parents know parents* has a negative and statistically significant effect on math-score gains, while *friends in school* has a positive and statistically significant effect. These findings support the hypothesis that horizon-expanding high schools foster more learning than do norm-enforcing high schools. When these two variables are interacted with the dummy variable for Catholic schools, one final conclusion is possible. The evidence for a negative net effect of *parents know parents* on mathematics learning comes predominantly from covariation in learning gains and *parents know parents* across public schools.

Does our article make too much of too little? This depends on sociologists' prior beliefs. We believe that most sociologists familiar with Coleman's theory maintain that the effect of parental social closure should be positive and that it should account for some substantial portion of the purported Catholic-school effect on learning. In reference to this prior belief, we view our findings as an important and novel contribution to the literature that hopefully will interest other researchers in the important but sometimes neglected subfield of sociology of education.

REFERENCES

Carbonaro, William J. 1998. "A Little Help from My Friend's Parents: Intergenerational Closure and Educational Outcomes." *Sociology of Education* 71:295–313.

———. 1999. "Opening the Debate on Closure and Schooling Outcomes (Comment on Morgan and Sørensen)." *American Sociological Review* 64:682–86.

Coleman, James S. 1990. *Foundations of Social Theory*. Cambridge, MA: Harvard University Press.

Hallinan, Maureen T. and Warren N. Kubitschek. 1999. "Conceptualizing and Measuring School Social Networks (Comment on Morgan and Sørensen)." *American Sociological Review* 64:687–93.

Morgan, Stephen L. 1998. "Social Capital, Capital Goods, and the Production of Learning." Paper presented at Social Capital: An International Conference, April 21, Michigan State University, East Lansing, Michigan.

Morgan, Stephen L. and Aage B. Sørensen. 1999. "Parental Networks, Social Closure, and Mathematics Learning: A Test of Coleman's Social Capital Explanation of School Effects." *American Sociological Review* 64:661–81.

Winship, Christopher. 1998. "Multicollinearity and Model Misspecification: A Bayesian Analysis." Presented at the winter meeting of the American Sociological Association's Methodology Section, April 5, Chicago, Illinois.