



## Discussion

## Estimating peer effects on health in social networks: A response to Cohen-Cole and Fletcher; and Trogdon, Nonnemaker, and Pais<sup>☆</sup>

J.H. Fowler<sup>a</sup>, N.A. Christakis<sup>b,c,\*</sup>

<sup>a</sup> Department of Political Science, University of California, San Diego, United States

<sup>b</sup> Department of Health Care Policy, Harvard Medical School, Boston, MA, United States

<sup>c</sup> Department of Sociology, Cambridge, MA, United States

### ARTICLE INFO

#### Article history:

Accepted 2 July 2008

Available online 9 August 2008

In July of 2007, we published a paper in the *New England Journal of Medicine* that used dynamic data over 32 years from the Framingham Heart Study Social Network (FHS-Net) to study the conditions under which obesity might spread from person to person (Christakis and Fowler, 2007, hereafter, CF). We found that obese persons formed clusters in the network at all time points and that these clusters extended to three degrees of separation (e.g., to a person's friend's friend's friend). Moreover, statistical analyses suggested that the clusters were not solely attributable to the selective formation of social ties among obese persons. A person's chances of becoming obese increased if he or she had a friend who became obese in a given time period.

Our analyses were restricted to adults, so a natural question to ask is whether or not the results would generalize to a population of adolescents. The existence of social norms regarding weight in both adults and adolescents should not be surprising, (Chang and Christakis, 2003) but, of course, causal inference in dyads, let alone in broader social networks, is difficult (Manski, 1993). These issues are addressed by two papers in this issue, "Is Obesity Contagious? Social Networks versus Environmental Factors in the Obesity Epidemic," by Ethan Cohen-Cole and Jason Fletcher (hereafter CCF) and "Peer Effects in Adolescent Overweight," by Justin Trogdon, James Nonnemaker, and Joanne Pais (hereafter TNP), and by a third working paper "Identifying Endogenous Peer Effects in the Spread of Obesity" by Timothy J. Halliday and Sally Kwak (Halliday and Kwak, 2007) (hereafter HK). All three of these papers analyze the same dataset and population, the National Longitudinal Study of Adolescent Health (Add Health), albeit with different methods and assumptions. Unlike the FHS-Net, which followed adults over a 32-year period (average age 38 in 1971), Add Health followed adolescents over a 7-year period (average age 16 in 1995).

All three sets of authors take the possibility of peer effects seriously, advancing the study of this important area. HK summarize their results by noting that they are able to replicate the pattern of results in our study, although their results are sensitive to specification of the dependent variable. If weight is characterized by a dichotomous variable indicating overweight (BMI > 25), then an association between friends is significant, but the association in the continuous measure of BMI is not. Similarly, TNP use a variety of econometric strategies to conclude that there are peer effects for obesity in the Add Health sample, especially among females and among adolescents with high BMI. For example, they use an

<sup>☆</sup> Further exchanges between the authors may appear on their home pages. Jason Fletcher <http://www.med.yale.edu/eph/faculty/fletcher.html>; Ethan Cohen-Cole <http://www.bos.frb.org/economic/econbios/cohen-cole.htm>; James Fowler <http://jhfwler.ucsd.edu/>; Nicholas Christakis <http://christakis.med.harvard.edu/>.

\* Corresponding author at: Department of Sociology, Harvard University, Cambridge, MA, United States.  
E-mail address: [christakis@hcp.med.harvard.edu](mailto:christakis@hcp.med.harvard.edu) (N.A. Christakis).

instrumental variables approach and a variety of definitions of endogenous peer groups to control for contextual effects. Here, we particularly address the CCF paper since it is the only one of the three papers that claims to reject the hypothesis that weight status can spread from person to person. To their credit, CCF exploit longitudinal data in a way that TNP and HK do not, but there are important problems in both their analysis and their interpretation of the results.

## 1. Getting the null right in CCF

In our view, CCF make a classic error in interpreting their own results. CCF seem to be single-mindedly focused on whether one can reject the null hypothesis of zero regarding whether weight gain in a social contact can cause weight gain in an ego. In three of their five specifications, they actually do replicate the approximate magnitude and the significance of the result from Framingham, which is quite remarkable. However, they interpret their remaining two results as a rejection of the Framingham result because they cannot reject the zero-null hypothesis.

What is crucial to note here, however, is that the Framingham estimates actually fall inside the confidence intervals of the CCF results for all five specifications. That is, CCF cannot reject the null hypothesis that the Framingham estimates are the true values. To put it even more starkly, there is a reasonable probability that the difference in the CCF and Framingham results is merely due to chance. Therefore, an alternative interpretation of their results is that they are actually supportive of the Framingham results (and also consonant with the conclusion reached in TNP and HK).

Moreover, as we will show below, the CCF results reflect a number of idiosyncratic and problematic econometric choices, and there are good empirical and theoretical reasons to be skeptical of their skepticism.

## 2. Contextual effects in the Add Health data

We attempted to replicate the CCF results but were unable to do so, even after several efforts to acquire information about specific details of CCF's modeling choices from CCF. We first relied on their manuscript and supplement, but it gives an incomplete description of the model; for example, it was unclear how income data were imputed. We then contacted CCF to ask for the code they used to generate their results. They were unwilling to share this code.

We therefore wrote new code to analyze the Add Health data to replicate their results; we sent it to them and asked them to comment. A brief response from CCF stated that a "difference in our approaches is that we 'lock in' Wave 1 friends—that is, we do not allow individuals to switch friends over time." This is an important and noteworthy omission from their description in the paper since Add Health (like the FHS-Net data) contains *dynamic* information about friendships at later waves, yet they chose instead to ignore this information and rely on a *static* representation of that data. This will cause them to assume some individuals continue to be friends when, in fact, they are not. This assumption by CCF stacks the deck against finding an effect, since it essentially adds "random" non-friend relationships (i.e., people who are no longer friends) to the pool of friends.

Since the CCF results cannot be replicated, and since they are in any case inconsistent with the conclusions of TNP and HK, we present our own results here based on the Add Health data, in an effort to figure out what is going on. CCF use only the first friend named by each subject, but when we interact the friend list order with alter's contemporaneous obesity, we cannot reject the hypothesis that order is irrelevant to the strength of the effect. We therefore use all observed friendships to maximize efficiency. CCF impute missing income and education data, but it is unclear how they do this. We use expectation maximization with importance sampling to impute all missing data (King et al., 2001). This procedure is widely used and well known to reduce bias relative to alternatives such as listwise deletion (Rubin, 2004; King et al., 2001). Like CCF, we also include variables for age, gender, race, ethnicity, and a fixed effect for the wave of the observation. And, as already noted, we use dynamic friendship data available from Waves 1–3 to update friendships at each wave, rather than making the assumption made by CCF that friendships are static and all people who became friends in junior high school retain all of their relationships into adulthood.

Table 1 presents the results. In model 1, notice that the coefficient of 0.033 on alter's contemporaneous obesity is significant ( $p=0.02$ ). In model 2, we add trends in weight change specific to each school as suggested by CCF. Not only does this addition have no effect on the induction effect, it also fails to achieve significance (CCF do not report the coefficient or standard error for their school trend effect, so we cannot compare results). However, what is clear is that the addition of school-specific trends does not matter in a model that uses all available data.

CCF appear dismissive of evidence that the induction effect in our Framingham data is *directional*. That is, if Mark names John as a friend, we expect John to have an effect on Mark. However, if John does not reciprocate by naming Mark as a friend, then John may not be affected by Mark's opinions or health behaviors. We hypothesize that influence flows from the named friend to the person who named him, but not necessarily vice versa, and in Framingham we find exactly that. The named friend has a significant effect on the namer but the namer has no significant effect on the named. New work in the econometrics of networks confirms that exploiting directionality in networks is a useful identification strategy (Bramoullé et al., 2007). Fortunately, Add Health also allows us to test the directional hypothesis. In model 3 of Table 1, we show that there is no evidence that named friends are influenced by namers ( $p=0.90$ ), confirming our results in the FHS-Net and providing additional evidence in favor of the causal interpretation regarding social influence in weight behaviors.

It is interesting to note that CCF did not report this result, since it speaks directly to their main point. If contextual effects are spuriously driving the relationship between ego and alter, then there is no reason to expect a directional result. The

**Table 1**  
Replication of Add Health results show school trends do not matter and show evidence of directional effects

	Model 1: alter is the named friend			Model 2: effect of named friend with school trends			Model 3: alter is the person who named the friend		
	Coef.	S.E.	<i>p</i>	Coef.	S.E.	<i>p</i>	Coef.	S.E.	<i>p</i>
Contemporaneous alter BMI	0.033	0.014	0.022	0.033	0.014	0.020	0.002	0.014	0.900
Alter previous BMI	−0.004	0.015	0.762	−0.005	0.015	0.759	0.002	0.015	0.908
Ego previous BMI	0.952	0.015	0.000	0.952	0.015	0.000	0.937	0.016	0.000
Ego gender	−0.129	0.085	0.129	−0.132	0.085	0.121	−0.176	0.100	0.078
Ego age	−0.079	0.030	0.008	−0.079	0.030	0.008	−0.071	0.033	0.033
Ego Hispanic	0.185	0.139	0.182	0.179	0.139	0.198	0.029	0.147	0.845
Ego Black	0.225	0.120	0.060	0.225	0.120	0.060	0.192	0.136	0.158
Wave	2.736	0.158	0.000	2.725	0.159	0.000	2.699	0.174	0.000
Maternal education	−0.041	0.020	0.037	−0.041	0.020	0.035	−0.023	0.021	0.279
Family income	0.000	0.001	0.642	0.000	0.001	0.652	−0.001	0.001	0.425
School trends				0.037	0.032	0.254			
Constant	−2.874	0.537	0.000	−2.854	0.540	0.000	−2.001	0.639	0.002
Deviance	55,734			55,732			27,782		
Null deviance	240,646			240,646			107,965		
<i>N</i>	8,504			8,504			4,742		

context should cause the named friend and the namer to move up and down simultaneously; hence, if we find a significant effect in one direction, we should also find it in the other: the namer should appear to have an influence on the named friend. Since we do not find such a significant effect, we believe the evidence in Add Health is suggestive of a causal effect, just as it is in Framingham. Obesity appears to spread from person to person.

CCF assert that our model specification in Framingham does not capture any contextual effects that vary across geographic space. It is important to note that friends in Add Health are all physically proximate (they are in the same school), whereas they are not in Framingham. If our estimates are biased because they capture community-level correlation, one implication is that the increased distance between friends will reduce the effect size (since distant social contacts are not contemporaneously affected by community-level variables). We specifically find that the relationship does *not* decay with physical distance, even up to hundreds of miles away (Fig. 3 in CF). We strongly emphasize this point in our article, and it was widely reported in the popular press, so it is difficult to understand how CCF could have missed this. The implication of this observation is that any contextual effects that are geographic in nature probably do not have an effect on the association between ego's and alter's obesity. In addition, as reported in CF, we also analyzed whether there was a relationship in weight behaviors between individuals and their immediate neighbors living at adjoining housing units. If contextual effects were driving the association in weight between social contacts, we would expect neighbors to appear to have an effect on each other, but we found that they do not.

Finally, in this regard, it is also worth noting that in Framingham we include dummy variables as controls for each exam, which effectively controls for average weight change in the whole population (time-specific effects). This might not capture regional variation within time blocks, but we also control for ego-specific factors that would account for some regional variation, including age, gender, and education.

### 3. Fixed effects in the Framingham heart study

Adding fixed effects to dynamic panel models with many subjects and few repeat observations creates severe bias towards zero coefficients. This has been demonstrated both analytically (Nickell, 1981) and through simulations (Nerlove, 1971) for OLS and other regression models and has been well known by social scientists, including economists, for a very long time. In fact, CCF even note that they do not add fixed effects to their logit regression model for this reason, but they strangely assert that fixed effects are necessary in the OLS model.

Since fixed effects generate bias towards zero coefficients, an especially conservative test of the Framingham results would be to subject them to this model. If we get a significant result in spite of the downward bias, then it would indicate our results are particularly robust. Table 2 shows that when we include fixed effects in our FHS-Net analyses, we find a significant influence of friend's BMI on ego's BMI, and the size of the effect is about the same at 0.05. Thus, our results in Framingham are robust to the CCF specification. Their assertion that a fixed effects model would change the Framingham result is thus not correct.

### 4. Homophily in networks

CCF imply that our model does not account for selection effects or homophily, which is the tendency of people with similar attributes to form ties (McPherson et al., 2001; Fowler, Dawes, and Christakis, 2008). If obese people befriend other obese people, it might create a correlation in weight status that is not driven by induction. This is a well-known issue in the analysis of such data (Carrington et al., 2005). CCF also assert that their method works in a way ours does not when they say it "accounts for self-selection of friends (homophily) according to weight status by looking only at the change in BMI from the time of declaration of friendship until the subsequent weight measurement." In fact, we did exactly the same thing as

**Table 2**  
New model of Framingham heart study social network data shows Framingham BMI result is robust to the inclusion of ego fixed effects

	Alter is the named friend		
	Coef.	S.E.	<i>p</i>
Contemporaneous alter BMI	0.053	0.018	0.003
Alter previous BMI	-0.016	0.019	0.387
Ego previous BMI	0.209	0.020	0.000
Ego age	0.134	0.075	0.073
Wave 3	-0.192	0.346	0.578
Wave 4	-0.308	0.603	0.610
Wave 5	-0.373	0.874	0.670
Wave 6	-0.588	1.174	0.617
Wave 7	-0.879	1.380	0.524
843 fixed effects not shown	-	-	-
Constant	14.547	10.567	0.169
<i>R</i> <sup>2</sup>	0.88		
<i>N</i>	2396		

CCF in our Framingham analysis, including only individuals who were friends at both times  $t$  and  $t + 1$ . In addition, we used generalized estimating procedures to account for individuals' repeated appearances in the data. And, unlike CCF, we tracked the friendship status of people across time.

In short, this model conditions on the initial weights of the two parties connected via a tie. Additional Monte Carlo simulation results documenting that homophily (ranging from no homophily to complete homophily) does not result in bias in the estimates of induction in this model specification are available in the full version of the paper at the authors' websites.

## 5. Inter-personal health effects

Social scientists have increasingly been exploring inter-personal health effects in the last few years (Burke and Heiland, 2006; Hammond and Epstein, 2007; Cutler and Glaeser, 2007; Clark and Loheac, 2007). And the extent of such effects may extend considerably beyond health behaviors such as obesity and smoking (Christakis and Fowler, 2007; Christakis and Fowler, 2008). For example, it is not hard to imagine that people emulate each other when it comes to symptoms as well. When a person's friends, co-workers, and family come to complain of depression, (Fowler and Christakis, in press) back pain, (Raspe et al., 2008) itching, cough, or headaches, they might do so as well.

Such peer effects are of obvious policy significance. First, it means that clinical and policy interventions may be more cost-effective than policy-makers may have previously supposed and that some interventions may gain more than others in the accounting (Christakis, 2004; Rosen, 1989; Powell et al., 2005; Harris and López-Valcárcel, 2008). Interventions that have greater positive externalities may rise in the analyst's estimation. If it costs \$25,000 to replace a man's hip and he gains four quality-adjusted life years from this intervention, and if his wife also gains one quality-adjusted life year as a result of having a more active partner, then the cost-effectiveness of the surgery has just gone up by 25%. But if a knee replacement does not benefit a spouse, then its cost-effectiveness does not rise. If we spend \$500 to get a person to quit smoking and if this person's quitting in turn results in one out of ten of her social contacts quitting, and if that leads to one out of *that* person's social contacts quitting as well, we can see that three people have quit for the price of one, tripling the cost-effectiveness of the intervention.

These kinds of effects are rarely taken into account by policy makers or even by entities with a collective perspective, such as insurers. Yet they probably should be.

A second implication of people's embeddedness in social networks is that group-level interventions may be more successful and more efficient than individual interventions. In fact, programs such as Alcoholics Anonymous, runners' clubs, symptom support groups (e.g., for chronic fatigue, breast cancer, eating disorders, psychiatric conditions), and weight loss groups are explicitly designed to create a set of artificial social network ties.

Finally, a social network perspective suggests that it may be possible to exploit variation in people's social network position to target interventions where they might be most effective in generating benefits for the group. For example, if funds are limited, it may be best to target people who are most likely to influence others. This has long been a focus in sociology (e.g., Valente and Pumpuang, 2007) and is increasingly becoming a focus in both health economics (e.g., Rao et al., 2007; Banerjee et al., 2007) and development economics.

People are interconnected, and so their health is interconnected.

## Acknowledgement

This work was supported in part by NIA P-01 AG-031093 and by a Pioneer Grant from the Robert Wood Johnson Foundation.

## References

- Banerjee, R., Cohen-Cole, E., Zanella, G., 2007. "Demonstration Effects in Preventive Care," Working Paper.
- Bramoullé, Y., Djebbari, H., Fortin, B., 2007. "Identification of Peer Effects through Social Networks," IZA Discussion Papers 2652, Institute for the Study of Labor (IZA).
- Burke, M.A., Heiland, F., 2006. "Social Dynamics of Obesity". FRB of Boston Public Policy Discussion Paper No. 06-5. Available at SSRN: <http://ssrn.com/abstract=934421>.
- Carrington, P.J., Scott, J., Wasserman, S., 2005. *Models and Methods in Social Network Analysis*. Cambridge University Press, Cambridge.
- Chang, V.W., Christakis, N.A., 2003. Self-perception of weight appropriateness in the U.S. *American Journal of Preventive Medicine* 24, 332–339.
- Christakis, N.A., 2004. Social networks and collateral health effects—have been ignored in medical care and clinical trials, but need to be studied. *British Medical Journal* 329, 184–185.
- Christakis, N.A., Fowler, J.H., 2007. The spread of obesity in a large social network over 32 years. *New England Journal of Medicine* 357, 370–379.
- Christakis, N.A., Fowler, J.H., 2008. The collective dynamics of smoking in a large social network. *New England Journal of Medicine* 358, 2249–2258.
- Clark, A.E., Loheac, Y., 2007. It wasn't me, it was them! Social influence in risky behavior by adolescents. *Journal of Health Economics* 26, 763–784.
- Cutler, D.M., Glaeser, E.L., 2007. "Social Interactions and Smoking." NBER Working Paper No. W13477.
- Fowler, J.H., Christakis, N.A., in press. "The Dynamic Spread of Happiness in a Large Social Network."
- Fowler, J.H., Dawes, C.T., and Christakis, N.A., 2008. A Model of Genetic Variation in Human Social Networks. arXiv: 0807.3089v1 (available at <http://arxiv.org/abs/0807.3089v1>).
- Halliday, T.J., Kwak, S., 2007. "Identifying Endogenous Peer Effects in the Spread of Obesity," Working Paper 7-27, University of Hawaii at Manoa, Department of Economics.
- Hammond, R.A., Epstein, J.M., 2007. "Exploring Price-Independent Mechanisms in the Obesity Epidemic." Center on Social and Economic Dynamics Working Paper No. 48. Available at SSRN: <http://ssrn.com/abstract=1024287>.
- Harris, J.E., López-Valcárcel, B.G., 2008. Asymmetric peer effects in the analysis of cigarette smoking among young people in the United States, 1992–1999. *Journal of Health Economics* 27 (2), 249–264.

- King, G., Honaker, J., Joseph, A., Scheve, K., 2001. Analyzing incomplete political science data: an alternative algorithm for multiple imputation. *American Political Science Review* 95 (1), 49–69.
- Manski, C., 1993. Identification of endogenous social effects: the reflection problem. *Review of Economic Studies* 60, 531–542.
- McPherson, M., Smith-Lovin, L., Cook, J.M., 2001. Birds of a feather: homophily in social networks. *Annual Review of Sociology* 27, 415–444.
- Nerlove, M., 1971. Further evidence on the estimation of dynamic economic relations from a time series of cross sections. *Econometrica* 39 (2), 359–387.
- Nickell, S., 1981. Biases in dynamic models with fixed effects. *Econometrica* 49 (6), 1417–1426.
- Powell, L.M., John, A.T., Hana, R., 2005. The importance of peer effects, cigarette prices and tobacco control policies for youth smoking behavior. *Journal of Health Economics* 24 (5), 950–968.
- Rao, N., Mobius, M.M., Rosenblat, T., 2007. Social Networks and Vaccination Decisions" (November 6, 2007). FRB of Boston Working Paper No. 07–12. Available at SSRN: <http://ssrn.com/abstract=1073143>.
- Raspe, H., Hueppe, A., Neuhauser, H., 2008. Back Pain: a communicable disease? *International Journal of Epidemiology* 37, 69–74.
- Rosen, M., 1989. On randomized controlled trials and lifestyle interventions. *International Journal of Epidemiology* 18, 993–994.
- Rubin, D., 2004. Multiple imputation for nonresponse in surveys. Wiley Press, New York.
- Valente, T.W., Pumpuang, P., 2007. Identifying opinion leaders to promote behavior change. *Health Education and Behavior* 34, 881–896.