

Detecting implausible social network effects in acne, height, and headaches: longitudinal analysis

Ethan Cohen-Cole, financial economist,¹ Jason M Fletcher, assistant professor²

¹Federal Reserve Bank of Boston, 600 Atlantic Avenue, Boston, MA 02210, USA

²Yale University, School of Public Health, 60 College Street, New Haven, CT 06510

Correspondence to: J Fletcher
jason.fletcher@yale.edu

Cite this as: *BMJ* 2008;337:a2533
doi:10.1136/bmj.a2533

ABSTRACT

Objective To investigate whether “network effects” can be detected for health outcomes that are unlikely to be subject to network phenomena.

Design Statistical analysis common in network studies, such as logistic regression analysis, controlled for own and friend’s lagged health status. Analyses controlled for environmental confounders.

Setting Subsamples of the National Longitudinal Study of Adolescent Health (Add Health).

Participants 4300 to 5400 male and female adolescents who nominated a friend in the dataset and who were both longitudinally surveyed.

Measurements Health outcomes, including headache severity, acne severity, and height self reported by respondents in 1994-5, 1995-6, and 2000-1.

Results Significant network effects were observed in the acquisition of acne, headaches, and height. A friend’s acne problems increased an individual’s odds of acne problems (odds ratio 1.62, 95% confidence interval 0.91 to 2.89). The likelihood that an individual had headaches also increased with the presence of a friend with headaches (1.47, 0.93 to 2.33); and an individual’s height increased by 20% of his or her friend’s height (0.18, 0.15 to 0.26). Each of these results was estimated by using standard methods found in several publications. After adjustment for environmental confounders, however, the results become uniformly smaller and insignificant.

Conclusions Researchers should be cautious in attributing correlations in health outcomes of close friends to social network effects, especially when environmental confounders are not adequately controlled for in the analysis.

INTRODUCTION

Providing credible estimates of the effects of social networks in choices and outcomes in health is important for increasing our understanding of the determinants of such outcomes as well as suggesting policies that could lead to improving health. Policies that are able to influence network effects could produce “social multipliers” in the benefits, where the total effect of the policy is magnified via social networks.

Recent work has examined potential social network effects in outcomes such as obesity, smoking, and alcohol use. Christakis and Fowler have presented

evidence of the person to person spread of obesity and quitting smoking among friends.^{1,2} Raspe et al proposed that back pain might be a “communicable disease.”³

Many methods used to estimate social network effects are subject to potentially large biases that result in the increased likelihood of detecting social network effects where none exists. For example, the use of standard econometric methods from literature on peer effects substantially reduces evidence of social network effects in obesity.⁴ Previous work that claimed to find social contagion in the diffusion of prescription drugs was confounded by marketing effects.⁵

Using standard methods we examined whether one can “find” network effects using common methods even in health outcomes that are unlikely to be transmitted socially: acne, headaches, and height.

METHODS

Empirical issues in estimating social network effects

The general form of the regression equation is typically some version of

$$health_{igt} = \delta health_{jgt} + \beta X_{it} + \mu_{gt} + \epsilon_{igt}$$

where a health outcome of individual i in reference group g at time t is determined by the health of other individuals in the reference group, own characteristics X , group level factors μ , and an error term. We focused on two main empirical difficulties in estimating social network effects (δ) within reference groups: firstly, that friendship selection is non-random, which leads to correlation between the error term and friend’s health, and, secondly, that confounding factors affect all members of the reference group (μ). There are many other potential econometric issues, including the reflection problem⁶—it is difficult to separate individual A’s effect on individual B from individual B’s effect on individual A—and the proper definition of the relevant reference group.

The first difficulty we focus on, called selection in the economics literature and homophily elsewhere, creates correlations in health outcomes because individuals in good (or bad) health tend to associate with other individuals in good (or bad) health. This non-random pattern of association across individuals can lead to correlations in health outcomes between friends that are not caused by direct social network effects.

The second difficulty, environmental confounding, can occur when a feature of the shared environment affects all individuals in the same reference group. For example, a fast food restaurant, convenience store, or gym opening near a school could simultaneously affect the weight of all individuals in networks within the school. Importantly, the presence of (often unmeasured) shared surroundings can lead to erroneously implicating social network effects in individual outcomes where none exists.

There are several common approaches to overcoming the difficulties outlined above.⁷ The problem of selection has been addressed through the use of fixed effects or random assignments of reference groups, such as the practice of randomly assigning first year roommates at some colleges.⁸ Other authors attempt to overcome selection effects by including lagged variables in the empirical model, though in general this practice produces biased results.⁴ Importantly, new research that combines multiple strategies to address the multiple difficulties is an important next step in detecting social network effects.⁹

The problem of confounding has generally been addressed by controlling for a rich set of individual, family, and environmental characteristics or using fixed effects at the group level.¹⁰ The problem, of course, is that social groups are often faced with similar environmental characteristics. If these are neglected, one can improperly interpret the results to imply that true “network effects” exist.

We argue that the test statistics drawn from what we call the “standard” approach are incorrect. In particular, because research has not accounted for the problems above, the standard errors from the simple models will be unreasonably small. Of course, if standard errors are too small, a researcher will be more likely to reject, incorrectly, any given null hypothesis.

It is simple in empirical work to assume that whatever information is available in the dataset is the same information that describes the social environment in which people live. In particular, it is common to assume that the set of variables are the appropriate ones to differentiate environmental confounders from true network effects. The problem is that the datasets used were rarely, if ever, constructed with this type of analysis in mind. As such, the individual and group characteristics are typically more appropriate for evaluating individual level health outcomes rather than group level interactions. To differentiate between network effects of obesity and confounders, for example, one would want to know the pattern of fast food restaurants or the caloric content of the school cafeteria available to the social network. Inclusion of the individual’s race, income, etc, might be reasonable proxies for some studies but cannot distinguish two otherwise similar groups that have different environments. Estimating a regression of any type without this salient information might show a “network effect” if one school is next to a fast food restaurant and another is not.

Though there are different approaches to estimating social network effects, we focused on the approach used in Christakis and Fowler¹²:

$$\text{health}_{igt} = \beta_1 \text{health}_{igt-1} + \beta_2 \text{health}_{jgt-1} + \beta_3 X_{it} + \varepsilon_{igt}$$

This addresses selection issues by controlling for lagged health outcomes. The claim is that “the use of a lagged independent variable for an alter’s weight status controlled for homophily.” (An alter is a person linked to the focal person—a friend.) Unfortunately, unless selection is conditioned only on this variable, this statement might be spurious. For example, if friendships are formed based on characteristics like self esteem, and if self esteem affects both current weight and future weight in differing ways, then adjustment for current peer weight status will not capture the self selection of friends based on self esteem that also affects future weight.¹¹ Another example could be that current smokers who are considering quitting might become friends with individuals who they suspect are likely to successfully quit smoking in the future. Controlling for

Table 1 | Summary statistics from National Longitudinal Study of Adolescent Health (Add Health) used to create sample for analyses; two observations per person

Variable (No of observations)	Mean (SD)	Range
Skin (n=4590)		
No with skin problem	0.05 (0.22)	0-1
No with friend with skin problem	0.05 (0.21)	0-1
Male	0.44 (0.50)	0-1
Age (years)	18.93 (3.12)	13-27
Black	0.17 (0.38)	0-1
Hispanic	0.15 (0.36)	0-1
Maternal education (years)	13.29 (2.31)	0-17
Family income (\$1000s)	48.77 (45.34)	0-900
Missing family information	0.29 (0.45)	0-1
Headache (n=5360)		
No with headache problem	0.06 (0.23)	0-1
No with friend with headache problem	0.05 (0.22)	0-1
Male	0.44 (0.50)	0-1
Age (years)	19.03 (3.15)	13-27
Black	0.18 (0.38)	0-1
Hispanic	0.15 (0.36)	0-1
Maternal education (years)	13.28 (2.29)	0-17
Family income (\$1000s)	48.18 (43.59)	0-900
Missing family information	0.29 (0.45)	0-1
Height* (n=4332)		
Height (m)	1.69 (0.10)	1.35-2.31
Height of friend (m)	1.68 (0.10)	1.37-2.11
Male	0.43 (0.50)	0-1
Age (years)	18.93 (3.12)	13-27
Black	0.17 (0.38)	0-1
Hispanic	0.15 (0.36)	0-1
Maternal education (years)	13.32 (2.31)	0-17
Family income (\$1000s)	49.05 (47.77)	0-900
Missing family information	0.28 (0.45)	0-1

*Originally measured in inches; 1 inch=2.54 cm.

Table 2 | Association between own health status and friend's health status (skin problems and headache) at baseline and extended specifications. Figures are odds ratios with robust standard errors

Outcome	Unadjusted	Adjusted*
Skin problems		
Control for environmental confounding	No	Yes
Lagged own skin problem	6.48† (1.09)	4.90† (0.85)
Friend's skin problem	1.62‡ (0.48)	1.23 (0.35)
Lagged friend's skin problem	0.89 (0.29)	0.75 (0.24)
Wave	1.78 (1.06)	1.56 (1.06)
Fixed effects	None	School
No of observations	4540	3856
Headache		
Control for environmental confounding	No	Yes
Lagged own headache problem	5.20† (0.86)	3.85† (0.62)
Friend's headache problem	1.47‡ (0.34)	1.14 (0.27)
Lagged friend's headache problem	1.13 (0.26)	1.09 (0.25)
Wave	0.67 (0.36)	0.77 (0.43)
Fixed effects	None	School
No of observations	5292	4750

*Sex, age, race, maternal education, family income, grade level in wave 1, indicator for missing family information.
†P<0.001.
‡P<0.1.

baseline smoking status of each individual might not be a sufficient control for self selection of friends. The general point is that individuals become friends for many reasons that are correlated with health and health trajectories, and adjustment for one measure of friend's health status is unlikely to control sufficiently for the non-random selection of friends. Finally, in the presence of social network effects, the use of lagged variables can lead to bias in estimation apart from the issues of self selection.¹²

The second issue with these methods—confounding—hinges on whether appropriate variables are included in the regression analysis. As discussed above, all students in schools that are near a fast food restaurant or that have limited access to exercise opportunities share an increased risk of obesity. As schoolmates are often also friends, these common environmental exposures might produce the appearance of “social network effects” if not controlled for in the empirical models, particularly if the type and direction of the friendship networks are used for adjustment. For example, if individual A declared himself a friend of individual B but not vice versa, then a social network effect should appear for A but not B. Unfortunately, while this “intransitivity” of the network can facilitate identification of network effects,¹³ it does not address confounding.

In addition to the equation above, we used logistic regressions analysis and generalised estimating equations to produce our baseline results for our binary dependent variables and ordinary least squares for our continuous dependent variable (we present the logistic and ordinary least squares results, but the generalised estimating equations method yielded nearly identical estimates).

Data source

We use the Add Health dataset to examine social network effects in three health outcomes for a national sample of adolescents. A full description of the sample design, data, and documentation is available at www.cpc.unc.edu/addhealth.

As we intended to investigate potential biases in previous methods, we looked at three health outcomes that could not credibly be subject to social network effects and were available in all three waves of the data: self reports of skin problems, self reports of headaches, and height over time. Among the many health outcomes available in the dataset, we found only three for which a social effect was truly implausible. In waves 1 and 2 of the survey individuals were asked, “How often in the past 12 months have you had skin problems, such as itching or pimples.” Answers included “never,” “just a few times,” “about one a week,” “almost every day,” and “every day.” In wave 3, they were asked whether they had taken prescription medication for acne in the past 12 months. To create comparable measures across waves, we created an indicator variable for waves 1 and 2 for whether the respondent reported having skin problems “almost every day” or “every day.” About 5% of the respondents reported this frequency, which was similar to the 5% of wave 3 respondents who reported having prescription medication for acne.

Likewise, in waves 1 and 2 of the survey individuals were asked, “How often have you had a headache in the last 12 months.” In wave 3, respondents were asked whether they had taken prescription medication for headaches in the last 12 months. To create comparable measures, we created an indicator variable for waves 1 and 2 for whether the respondent reported a headache “almost every day” or “every day.” Five percent of the respondents reported this frequency, which was similar to the 5% of wave 3 respondents who reported having prescription medication for headaches. Our inclusion of a wave dummy variable should control for the difference in headache questions across waves.

For our height outcome, we use self reported height across all three waves, which was reported in feet and inches. Height was measured in waves 2 and 3 of the survey but self reported in wave 1. We made several corrections to the self reported height variables. We did not allow individuals to lose height over time. We also coded heights that increased more than 8 inches (20 cm) in one year (between wave 1 and 2) as errors. Our results were quantitatively similar whether or not we controlled for this measurement error in the height variable.

We had information on friends for over 5000 individuals, about 2000-3000 of whom were followed over time along with at least one same sex friend, depending on the health outcome. For individuals for whom health information on more than one friend was available, we selected the friend with the highest nomination (1st-5th). Nearly two thirds of the individuals in our sample were matched to only one friend's health because of the sample design. We selected only

one friend to be consistent with previous research, in which the respondents had on average 0.7 nominated friends. These sample sizes gave us about 4000 person year observations for each analysis. Table 1 shows summary statistics for our three samples.

Though there are several important differences between the Add Health and the Framingham Heart Study used in previous research,¹² the two datasets are sufficiently similar to use to evaluate the role of transmission mechanisms.⁴ In particular, previous work has shown the ability to largely replicate the social network effects for obesity with the methods used by Christakis and Fowler¹ and the Add Health data.⁴

RESULTS

Table 2 shows baseline estimates for self reported skin problems, with unadjusted data and data adjusted for environmental confounding through the use of school fixed effects (see below). Our baseline unadjusted results suggest that having a friend with skin problems increases the respondent's chances of skin problems (odds ratio 1.62, 95% confidence interval 0.91 to 2.89). While the unadjusted result is significant only at the 10% level, our intention is not to state that the effects are relevant but to show that the findings are fragile. The 5% standard, typically drawn from Fisher's recommended threshold, was never intended to be a strict one, but rather reasonable guidance as a threshold for rejection of the null.¹⁴

This empirical model generates such large results that we would reject a null at the 10% level even when the true contagion effect is zero. The magnitude of our result (relative risk of 1.58) is similar to the 57% increase in risk of becoming obese when a friend is obese, as reported elsewhere,¹ and larger than the 36% increase in quitting smoking when a friend quits.²

Table 2 also shows the results for self reported headache problems. Our unadjusted baseline results suggest that having a friend with headache problems increases the respondent's chances of headache problems (1.47, 0.93 to 2.33). Again, this result is marginally significant but quite large in magnitude. It is

larger than the reported social network effect for smoking² and nearly as large as the result for obesity.¹

In the adjusted results in table 2, we show that adding simple controls for environmental confounding reduces the "social network effect" by over 50% and renders the results indistinguishable from zero. We controlled for school level fixed effects in our empirical models to control for all environmental conditions shared by students in the same school. These results support theoretical work in econometrics that show that not controlling for confounders can lead to inflated estimates of "social network effects."⁶ Previous research controlled only for a limited set of individual level variables and time effects but no other shared environmental factors. Inclusion of school level fixed effects does not provide precise information on network specific confounders, but even a relatively blunt measure is effective at correcting spurious results. Our results suggest that failure to control for confounders increases the chances of attributing similarity in health outcomes between friends to social network effects when the similarities are caused by shared environments.

Finally, we examined the "social network effects" of height between friends. We used ordinary least squares regression analysis because of the continuous nature of the outcome and table 3 shows the results. Our baseline findings in column 1 suggest strong contagion in height between friends over time. This finding for height is driven by a different specification error that is particularly acute for height but might be more generally applicable to other health outcomes that do not change often over time. For the case of height, the correlation between waves is greater than 0.95. This high multicollinearity probably generates the "social network effects" in a case (height) where we would expect no true contagious effects. The association between friend's current height and the respondent's current height was 0.18, but much of this correlation can be explained with adjustment for lagged respondent height, but friend's height is still associated with own height with a P value < 0.10 and a small magnitude. Finally, we controlled for school fixed effects, which reduced the magnitude of the "social network effect" of height. We believe that the observed correlation in height across friendship declarations is due to selection.

DISCUSSION

The methods of detecting "social network effects" of health outcomes commonly found in the recent medical literature might produce effects where none exists. The presence of network effects in three health outcomes—headaches, skin problems, and height—disappeared after we controlled for environmental confounders. These methods might produce fragile results and consequently can produce premature claims of social network effects in health outcomes. Lack of controlling for confounding factors is not a solution in itself, but any individual study needs to fully articulate the necessary assumptions and explain how common identification issues apply to the study. In the

Table 3 | Association between own height and friend's height difficulties with multicollinearity and confounding*. Figures are odds ratios with robust standard errors

Variable	Col 1	Col 2	Col 3	Col 4
Lagged own height	0.89† (0.01)	—	0.88† (0.01)	0.881† (0.01)
Friend's height	0.21† (0.03)	0.18† (0.022)	0.01‡ (0.01)	0.009 (0.01)
Lagged friend's height	-0.21† (0.03)	—	—	—
Wave	1.37† (0.08)	0.11 (0.21)	1.678† (0.09)	1.580† (0.11)
Fixed effects				
Observations	4284	4284	4284	4284
R ²	0.91	0.53	0.91	0.91

*Additional controls: sex, age, race, maternal education, family income, grade level in wave 1, indicator for missing family information. Column 1 recreates previous specifications by controlling both own and friend's lagged height. Column 2 repeats column 1 except lagged own height and lagged friend's height are not controlled. Column 3 repeats column 2 with the addition of lagged own height. Column 4 repeats column 3 with the addition of controls for environmental confounding in the form of school level fixed effects.

†P<0.001.

‡P<0.1.

WHAT IS ALREADY KNOWN ON THIS TOPIC

Recent research has shown that individuals who are socially connected also engage in similar health behaviours and have similar health outcomes

Socially connected individuals might have similar outcomes because they share similar environments, because they purposefully select their connections, or because their connections causally influence their behaviours and outcomes

These competing hypotheses are difficult to distinguish using many current empirical models, which might lead to the detection of causal “social network effects” where none exists

WHAT THIS STUDY ADDS

Current empirical methods used to estimate causal social network effects might detect implausible network effects, including “contagion” in headaches, skin problems, and height between adolescent friends

Caution is needed in attributing causality in empirical studies of social network effects; empirical models are needed that can distinguish causal and non-causal channels of social influence

cases here, the absence of confounding factors is the most salient, though others might emerge in different contexts.

Strengths and weaknesses

The core strength of our study is that we used common empirical methods and sets of control variables to show that the evidence of social network effects can largely be eliminated after adjustment for environmental confounders. Indeed, given how fragile the results are in many cases, one may suspect that other tests of sensitivity would generate similar changes in significance. Weaknesses of the study include the marginal levels of significance for two of our health outcomes as well as data limitations with our ability to test additional implausible health outcomes within this sample. For example, it would have been ideal to include other physical characteristics such as foot size or head circumference. In two of our cases, we have P values larger than 0.05, a standard level by which many researchers claim significance, and smaller than 0.10. Though on their own we would not submit to the strength of these results, we believe that the change from a P value <0.10 to a much larger value supports two points. Firstly, even truly implausible effects can generate results that support the hypothesis, and, secondly, appropriate inclusion of confounders returns the error bounds to a failure to reject the null.

There is a need for caution when attributing causality to correlations in health outcomes between friends using non-experimental data. Confounding is only one

of many empirical challenges to estimating social network effects, but researchers do need to attempt to minimise its impact. Thus, while it will probably not be harmful for policy makers and clinicians to attempt to use social networks to spread the benefits of health interventions and information, the current evidence is not yet strong enough to suggest clear evidence based recommendations. There are many unanswered questions and avenues for future research, including use of more robust empirical methods to assess social network effects, crafting and implementing additional empirical solutions to the many difficulties with this research, and further understanding of how social networks are formed and operate.

We thank David Paltiel and the reviewers and editorial board for helpful comments that substantially improved the paper.

Contributors: JMF acquired and analysed the data and is guarantor. Both authors interpreted the findings and drafted and revised the manuscript.

Funding: No external funding was used to support this research.

Competing interests: None declared.

Ethical approval: The Yale HIC exempted this research.

Provenance and peer review: Not commissioned; externally peer reviewed.

- 1 Christakis N, Fowler J. The spread of obesity in a large social network over 32 years. *N Engl J Med* 2007;357:370-9.
- 2 Christakis N, Fowler J. The collective dynamics of smoking in a large social network. *N Engl J Med* 2008;358:249-58.
- 3 Raspe H, Hueppe A, Neuhauser H. Back pain, a communicable disease? *Int J Epidemiol* 2008;37:69-74.
- 4 Cohen-Cole E, Fletcher JM. Is obesity contagious? Social network vs. environmental factors in the obesity epidemic. *J Health Econ* 2008;27:1382-7.
- 5 Van den Bulte C, Lilien GL. Medical innovation revisited: social contagion versus marketing effort. *Am J Sociol* 2001;106:1409-35.
- 6 Manski C. Identification of endogenous social effects: the reflection problem. *Rev Econ Stud* 1993;60:531-42.
- 7 Durlauf SN. Neighborhood effects. In: Henderson JV, Thisse JF, Eds. *Handbook of regional and urban economics*. Vol 4. 1st ed. Amsterdam: Elsevier, 2004:2173-242.
- 8 Sacerdote B. Peer effects with random assignments: results from Dartmouth roommates. *Q J Econ* 2001;116:2.
- 9 Fletcher JM. *Social interactions and smoking: an IV/FE approach*. Yale University Working Paper. 2008. http://search.ssm.com/sol3/papers.cfm?abstract_id=1069912
- 10 Arcidiacono P, Nicholson S. Peer effects in medical school. *J Public Econ* 2005;89:327-50.
- 11 Wertheim EH, Paxton SJ, Maude D, Szmukler GI, Gibbons K, Hiller L. Psychosocial predictors of weight loss behaviors and binge eating in adolescent girls and boys. *Int J Eat Disord* 1992;12:151-60.
- 12 Liu X, Lee LF, Kagel J. *Dynamic discrete choice models with lagged social interactions: with an application to a signaling game experiment*. Ohio State: Mimeo, 2006.
- 13 Bramoullé Y, Djebbari H, Fortin B. *Identification of peer effects through social networks*. Bonn, Germany: Institute for the Study of Labor (IZA), 2007. (IZA Discussion Papers 2652.)
- 14 Sterne JAC, Davey Smith G, Cox DR. Sifting the evidence—what’s wrong with significance tests? *BMJ* 2001;322:226-31.

Accepted: 3 November 2008