

Incorporating social network effects into cost-effectiveness analysis: a methodological contribution with application to obesity prevention

Chad Konchak,¹ Kislaya Prasad²

To cite: Konchak C, Prasad K. Incorporating social network effects into cost-effectiveness analysis: a methodological contribution with application to obesity prevention. *BMJ Open* 2012;**2**:e001038. doi:10.1136/bmjopen-2012-001038

► Prepublication history and additional material for this paper are available online. To view these files please visit the journal online (<http://dx.doi.org/10.1136/bmjopen-2012-001038>).

Received 17 February 2012
Accepted 26 September 2012

This final article is available for use under the terms of the Creative Commons Attribution Non-Commercial 2.0 Licence; see <http://bmjopen.bmj.com>

¹Clinical Analytics Department, NorthShore University Health System, Evanston, Illinois, USA

²Robert H. Smith School of Business, University of Maryland, Maryland, USA

Correspondence to
Dr Kislaya Prasad;
kprasad@umd.edu

ABSTRACT

Objectives: To develop a methodology for integrating social networks into traditional cost-effectiveness analysis (CEA) studies. This will facilitate the economic evaluation of treatment policies in settings where health outcomes are subject to social influence.

Design: This is a simulation study based on a Markov model. The lifetime health histories of a cohort are simulated, and health outcomes compared, under alternative treatment policies. Transition probabilities depend on the health of others with whom there are shared social ties.

Setting: The methodology developed is shown to be applicable in any healthcare setting where social ties affect health outcomes. The example of obesity prevention is used for illustration under the assumption that weight changes are subject to social influence.

Main outcome measures: Incremental cost-effectiveness ratio (ICER).

Results: When social influence increases, treatment policies become more cost effective (have lower ICERs). The policy of only treating individuals who span multiple networks can be more cost effective than the policy of treating everyone. This occurs when the network is more fragmented.

Conclusions: (1) When network effects are accounted for, they result in very different values of incremental cost-effectiveness ratios (ICERs). (2) Treatment policies can be devised to take network structure into account. The integration makes it feasible to conduct a cost-benefit evaluation of such policies.

INTRODUCTION

Human beings are social creatures, embedded within social networks through which information and behaviours percolate. A number of influential recent studies have shown that these networks have the potential to impact an individual's health behaviours, decisions and outcomes. There is now considerable evidence that the existence of social support and social ties affects health.^{1–3} It has

ARTICLE SUMMARY

Article focus

■ A number of recent studies have shown that social networks can have a significant impact on an individual's health behaviours, decisions and outcomes. However, network effects are typically neglected in cost-effectiveness analysis (CEA). We develop a methodology for integrating social influences into CEA and illustrate its effectiveness with the example of obesity prevention.

Key messages

■ The presence of network effects can lead to very different values of incremental cost-effectiveness ratios.
■ Cost-effective treatment policies can be devised by taking network structure into account.

Strengths and limitations of this study

■ A novel model which is simple and easily implemented in a variety of scenarios.
■ For the obesity example, the properties of the network and estimates of parameters for the simulation are approximations and not drawn from a longitudinal study. Hence, results are to be interpreted as suggestive.

been found, for instance, that social isolation is related to increases in death rates.⁴ While it is well known that pathogens can spread through person-to-person contacts, recent studies point to the fact that health behaviours can also spread through social contact between networked individuals. Links between individuals create pathways for the transmission of influence, information and resources. Consequently, these links determine extant norms of behaviour and, ultimately, health outcomes in the population. One of the key insights of research in social networks has been that a person's location in a network can explain his or her behaviour, and variation in the structure of networks can explain population-wide outcomes.⁵ In the

field of public health, the value of network approaches in stopping the spread of infectious diseases is well established.^{6 7} In such contexts, effective containment strategies need to take account of properties of networks. With the growing realisation that a much broader set of health outcomes depend on properties of social networks, there is a need to account for network effects in the economic evaluation of public health strategies, and to utilise network properties to design effective approaches to managing diseases. This paper seeks to provide a framework for this using the example of a major public health challenge—the spread of obesity.

Several recent papers have explored the link between social network structure and patterns of obesity. Perhaps the most influential study in this literature is Christakis and Fowler,⁸ which was characterised in the popular media as finding that obesity is contagious. Their principal finding was that the probability of becoming obese increases when people with whom there are shared social ties become obese. One of the key points made by these authors is that the presence of social effects raises the possibility that the same powers that spread obesity could be harnessed to spread healthy behaviours to fight obesity. Since this research was published, many other papers have studied the relationship between social influence and obesity, seeking to identify the pathways through which social influence may affect weight. In a review of recent literature, Hammond⁹ points to three types of influence—social norms, social capital and social stress. Two types of norms appear to be important—norms of body type, and eating norms. Burke *et al*¹⁰ examine individuals' self-classification of weight (ie, whether overweight) compared to their actual BMI and find that even as the US population has become more overweight, individuals' self-classification (at a given BMI) has moved in the opposite direction. This is consistent with self-classification being socially determined. Other studies of self-perception and social comparison are consistent with this.¹¹ In a study of eating norms of 9-year-olds to 15-year-olds, Salvy *et al*¹² found that overweight youth tended to eat more when in the company of others who were also overweight. In another study that looked at the power of networks for combating obesity, Bahr¹³ found that an individual who is entrenched inside a network of overweight individuals (as opposed to being networked with normal weight individuals as well) would experience only temporary success with any weight loss programme because the power of the social network would dominate the individual, reversing any weight loss that had occurred. This property of social networks suggests that, when applying policies or strategies for abating obesity, simply targeting individuals may prove to be wasteful and inefficient unless the network the individual is involved in is simultaneously considered as well.

Social networks do not just play a role in the obesity epidemic. In a separate study targeting smoking cessation, Christakis and Fowler found that while the number

of smokers in the USA has fallen from 45% to 21% in the last four decades, the percentage of smokers in a given social network remains the same size, which suggests that groups of smokers are quitting together. This would indicate that their social networks play a key role in smoking cessation.¹⁴ In addition to smoking, peer effects are believed to be important for a range of other risky behaviours, especially among adolescents. As a final example, some choices, such as vaccination of babies, depend on information acquired from trusted sources, which may be other members of a network.

Despite the growing evidence in support of the importance of social ties for health, social influence and properties of networks have not been incorporated in cost-effectiveness analysis (CEA). In the typical CEA, some new treatment is compared against a baseline in terms of its costs and effectiveness (for comprehensive descriptions of the CEA methodology, see Gold *et al*¹⁵ and Drummond *et al*,¹⁶ and for examples of more recent applications, see Chan *et al*,¹⁷ Choudhury *et al*¹⁸ and Post *et al*¹⁹). The most common measure of effectiveness is quality-adjusted life years (QALYs), whereas costs are measured in dollars. The cost effectiveness of the new treatment is measured by the *incremental cost-effectiveness ratio* (ICER)—the ratio of incremental cost of the new treatment to its incremental effectiveness. It is useful, for the present purposes, to think of the ICER computation as follows. We start with a cohort of individuals, and imagine simulating their health histories. As time progresses, people may move from one health state to another with specified probabilities and they may die. Under different treatments, the amount of time spent in the various health states will differ. For instance, people may survive longer (on average) under a new public health regime which mandates annual screening for certain diseases. The average cost incurred over a lifetime will also differ across treatments. By simulating the life history of the cohort under different treatment regimes, we are able to calculate the incremental cost of every life year gained. Adjustments can be made for quality of life. For instance, we may judge every year spent living with obesity (and its attendant diseases) as worth less than a year in perfect health (equal to 11 months living in perfect health, say). In which case, the number of years lived can be scaled down appropriately. The important point here is that the simulation described above ignores social ties, if any, between members of the cohort. As far as we are aware, available tools for CEA do not allow for the possibility of such ties, and most published studies do not include them in their calculation.

In this paper, we propose a new methodology for CEA in the presence of network effects. Additionally, using obesity control as our test case, we construct a prototype of a CEA model that accounts for social networks as an influence on weight fluctuations. Our approach, in essence, is to run the simulation described above, but model social ties between members of the cohort. For

instance, each individual could belong to a small number of groups. The transition between states, for this individual, can then be made to depend upon the health states of individuals in these groups. For the rest, we can make the ICER computation exactly as before. In the specific context of policies to combat obesity, a number of CEA studies have been conducted and are summarised in Cawley.²⁰ The studies listed there do not allow for social influence, although one study²¹ does consider an intervention that takes the form of social support to promote walking. In our model, the presence of social ties influences the transition between weight categories. For instance, an individual is more likely to gain weight—and less likely to lose weight—if he or she has more ties with others who are obese. Since the mortality and morbidity associated with obesity are assumed to be higher, the nature of social ties has real effects on the number of years lived and the costs incurred over this time. Within this setting, we evaluate the effects of educating individuals on a diet plan. Treatments are assumed to be costly, and we compare two treatments. The first is to treat everyone who is overweight, while in the second we treat only those overweight individuals who are *boundary spanners* (belong to more than one group instead of being entrenched only within a group). We consider in particular the effects of changing the strength of social influence, and using network positions to target treatment. While there are many subtle ramifications, some of them depending on the specifics of our model, the principal inferences that emerge are that (1) social influence can have important effects on the cost effectiveness of treatments and (2) using our model, optimal treatment policies can be devised taking network structure into account.

Understanding how networks influence diseases will have significant implications for the treatment of diseases for which social interactions are important. We show that this understanding is likely to affect the cost-effectiveness trade-off, and may affect our determination of the optimal policy. We make novel contributions to the social network literature by integrating it with a Markov model of disease progression. At the same time, we contribute to the literature on CEA using Markov models by incorporating network structures and their effect on disease progression. We find that the method is easy to apply, and results are clear and easy to interpret.

METHODOLOGY

CEA involves simulating life histories of cohorts under alternative treatment policies. We describe here our approach to CEA. In our simulation, individuals are embedded in networks and behaviours are subject to social influence. Consequently, the life histories of individuals (especially their health) are subject to social influence. Our simulation model was developed using an object-oriented methodology programmed in Visual

Basic .NET, which implemented a web interface as the presentation tier. All data for the model was generated during runtime and processed by the application throughout the duration of the simulation.

Initialisation of the cohort

In the Christakis and Fowler⁸ study, the choice of social network was driven by data availability (they construct the social network of Framingham subjects using contact information that was recorded on tracking sheets to facilitate follow-up). Other studies, such as in the setting of a classroom, use surveys in which individuals are asked to list their friends and/or class leaders in their perception.²² In this paper, we construct a hypothetical network at random. All individuals belong to at least one group (which is, in graph-theoretic terms, a clique). Some individuals belong to more cliques than others. The groups are allowed to vary in size. An example of such a random network is displayed in figure 1, with nodes representing individuals (for this illustration the number of individuals is assumed to be 100, and the size of a node is proportional to the number of ties for the individual).

In our simulation study, the cohort comprised of a population of 100 000 individuals (of age 20) who, in the baseline model, were randomly placed in at least one and up to three groups. Individuals who are in more than one group will be called *boundary spanners*. They also tend to have a larger number of ties (higher degree centrality). Individuals are subject to direct social influence only from people who belong to the groups in which they are members. Ties will be assumed to be reciprocal, so that people exert influence on all individuals in groups to which they belong. Further, we do not allow for heterogeneity in types of ties. At the initial stage of the simulation the population was distributed into the following weight categories, which were also randomly assigned: 37% normal weight, 32% overweight and 31% obese. As a consequence, our virtual population consisted of 20 000 groups with an average of 10 members each and consisting of randomly distributed weights so that a group could initially have any distribution of normal, overweight and obese members. Figure 1 displays an initial state. We allow individual weight status to change. The rule governing this change is described next.

Transitions between states

At the end of each stage of the simulation, members in the population move between states (eg, from normal to overweight, normal to dead, normal to normal, etc) according to specified transition probabilities. Our approach is similar to that of traditional Markov models (see figure 2). However, probabilities determining weight changes are calculated dynamically for each individual based on information about their position in the network. The transition probability is computed from a predefined baseline probability, which is multiplied by




Figure 1 A sample social network (the number of individuals is 100). Blue indicates normal weight, yellow indicates overweight and red indicates obese. The network is graphed using UCInet 6 (Borgatti *et al*, 2002^{22a}).

an *influence factor* to account for the effect of social influence on weight change. The baseline probabilities (see figure 3) can be understood as the probability of state transitions absent any social influence. It is important to note that while these probabilities are not grounded in research on weight changes, being chosen for illustrative purposes only, they are not out of the realm of possibility for at least some demographic groups. In the simulation (without social influence), the first 10 year average growth rate of obesity is around 2% after which there is a levelling off of obesity prevalence at about 39%. At least to an approximation, this resembles recent US experience. For comparison, and calibrating from a more complicated dynamic process, Hill *et al*²³ report

for the Framingham data: ‘We find that the current rate of becoming obese is 2% per year and increases by 0.5 percentage points for each obese social contact. The rate of recovering from obesity is 4% per year, and does not depend on the number of non-obese contacts’.

Influence

When two individuals belong to the same group they are said to share a *primary connection*. An individual’s social network is assumed to be the set of individuals with whom they share a primary connection (in any of the one to three groups to which the individual belongs). The weight categories of these primary connections were aggregated to determine the influence on the individual of interest. If less than 50% of the connections were normal, then there was an increased tendency to gain weight (the probability to gain weight increased). If more than 50% of the connections were obese, then this tendency was made even stronger. An individual with an influence the same as their own (eg, normal weight with normal influence) had no change to their baseline probability. An individual with an influence one degree different than their own (eg, normal weight with overweight influence) had their baseline probability multiplied by the *influence factor* in the direction of the influence (eg, for an *influence factor* of 2, an overweight with an obese

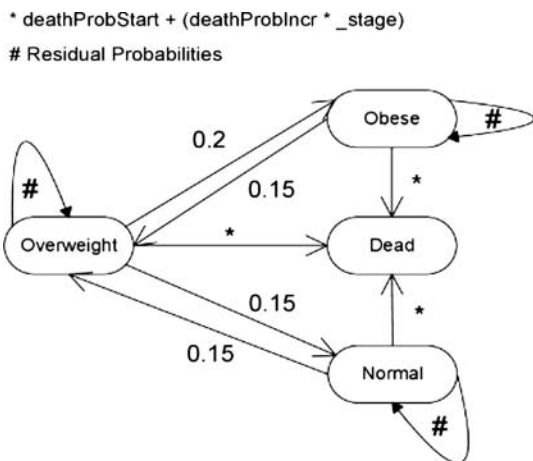


Figure 2 Markov diagram (# denotes probability of remaining in the same state; * is the age-adjusted probability of death from the given state.

	Normal Gain	Overweight Gain	Overweight Lose	Obese Lose
W/o Treat	0.15	0.2	0.15	0.15
W/ Treat	N/A	0.1	0.25	0.25

Figure 3 Probability of weight changes.

influence was twice as likely to gain weight, but half as likely to lose weight compared to the baseline). An individual with an influence two degrees in the direction of their own (normal weight with obese influence) had their baseline probability multiplied by the *influence factor* plus 1. The influence factor is a variable that can be input by the user as a parameter when the simulation is run. It can hold the value of 1 (for no influence), 1.5 (medium influence) or 2 (strong influence).

These numbers are chosen on the basis of prior research. For instance, Christakis and Fowler⁸ report: “If an ego stated that an alter was his or her friend, the ego’s chances of becoming obese appeared to increase by 57% if the alter became obese ... between mutual friends, the ego’s risk of obesity increased by 171%”. While there is some disagreement about the magnitude of social influence²⁴ these numbers appear reasonable, at least to an approximation.

Age effects

The probability of death is based on age. An iteration of the simulation is assumed to represent 1 calendar-year. A simple linear equation was used to increase the probability of death at each stage of the simulation. We begin the simulation with a cohort of age 20, where each individual is in one of three weight groups (normal, overweight or obese). At stage 0, the probability of death was 0.001 for all groups. Normal weight and overweight individuals shared the same probability of death at each stage, which increased by a value of 0.0002 per period, while obese individuals experienced an increase of 0.0006 per period. Note that the probability depends on age and weight status, but not weight history (hence, after 10 years of the simulation, if a person is obese his death probability would be $0.001 + 0.0006 \times 10$ regardless of whether they just became obese or had always been obese). In determining death probabilities we use information from the social security administration death actuarial tables (<http://www.ssa.gov/oact/STATS/table4c6.html>) assuming that we are starting at age 20 for men (and continuing for 80 years). We choose numbers for death probabilities for normal/overweight and obese populations such that there is a higher chance of death for the obese, but the average for the entire population corresponds to the expected death probabilities in the tables.

Calculating costs and effectiveness

At every stage of the simulation, individuals are in the different health states. We associate with every state a dollar value. In particular, at every stage of the simulation, a dollar value was applied to every individual according to their weight category (using the numbers in figure 4). These costs are taken from a CDC *Weight of the Nation* press briefing from Dr Eric A Finkelstein: “... normally an individual would spend \$3400 per year in medical expenditures and that number rises to about \$4870 if that individual is obese.” These numbers are based on a

Category	Cost
Living a year in normal weight	\$3400
Living a year overweight	\$3400
Living a year obese	\$4870
Following a normal diet	\$2830
Following a diet plan	\$4461

Figure 4 Costs.

published study^{25 26} These costs represent a conservative approach to estimating obesity by only considering medical expenditures and not taking into account any indirect expenses such as productivity costs, transportation costs and human capital costs, which are detailed in Hammond and Levine.²⁷ Average costs were calculated according to the proportion of individuals in each weight class (after excluding those dead). The cumulative cost was aggregated for each stage until stage 80. We measure effectiveness as simply the number of years lived. Since obesity is associated with a number of chronic diseases, it can be argued that a year lived in the obese state should be adjusted for quality (see, for instance, Maheswaran *et al*²⁸). The omission means that our results are conservative, understating the effectiveness of policies that are better in treating obesity. Effectiveness was also calculated at each stage by taking the proportion of those not dead and aggregating those data for each stage.

Treatments

Treatment takes the shape of a diet plan. We ran our simulation through three treatment plans: No Diet Plan, Treat All Diet Plan and Treat Boundary Spanners Diet Plan. The first two plans are self-descriptive. In the third plan, we attempt to utilise individuals’ positions in the social network in the design of the treatment plan. Based on the conjecture that individuals (nodes) who span boundaries have a larger influence on population-wide outcomes, we restrict treatment to only individuals in more than one group. An individual on a diet plan had a lower probability of gaining weight and higher probability of losing weight (see figure 3). For each stage, an individual not on a diet plan was assessed a normal diet cost, while an individual on the diet plan was assessed a higher diet cost for the stage (figure 4). Diet plan costs were taken from the *Forbes Magazine* article ‘Costly Calories’ that analysed weekly sample menus from 10 of the most popular diet plans from official publications and found the median weekly cost to be \$85.79 per week. This same publication also noted the average single American spends \$54.44 per week on food.²⁹ Only overweight and obese individuals were eligible for the diet plan. At any stage, in the *Treat All* plan, all overweight and obese individuals are treated. In the *Treat Boundary Spanners* plan only a subset of these individuals are treated, based on their positions within the

network. Once these three treatment plans were applied, we were able to determine the ICERs for both the Treat All plan and the Treat Boundary Spanners plan as compared with the baseline treatment, which is Treat None.

Testing the model

We ran our simulation through the same three treatments three separate times with a different *influence factor* value for each run. The values used were 1, 1.5 and 2. Setting the influence factor to 1 (no social networking) we were able to test our model to ensure that the probabilities, costs and effectiveness were being calculated correctly. We implemented a traditional cost-effectiveness Markov model using the widely used *TreeAge* decision analysis software. Running both the *TreeAge* model with the same values as our custom simulation (with the *influence factor* set to 1), we were able to generate identical numbers. As a consequence, this first series became our baseline series with which to compare social networks. Finally, we ran the simulation through two more series with the *influence factor* set to 1.5 and 2 in order to develop sets of result with which to analyse the effects of social influence on cost effectiveness. Additionally, a range of sensitivity analyses (to be described below) were performed to explore dependence of ICERs on costs, efficacy of treatments and—in particular—network structure.

Sensitivity analysis

As part of the sensitivity analysis (of the baseline case), we considered:

1. The effect of variations in costs of living with obesity.
2. The effects of variations in the effectiveness of treatments. In other words, how is cost effectiveness affected by variations in the probability of weight gain and loss with treatment? For this, we examined the effects of scaling the probabilities with treatment (in figure 3) up or down.
3. Relaxations of the assumption that groups are randomly assigned individuals with different weights. Since there is evidence that obesity differs according to socio-economic status, race and other criteria, it is important to know what happens when, initially, there is segregation according to weight types.
4. The effect of changes in network structure.

RESULTS

Our main results are presented in figure 5. The column on costs shows the average accumulated lifetime costs (including both the cost of living out a particular health history, as well as the costs of the diet) under the different treatments. As expected, treatment adds to costs. The effectiveness column shows the expected number of years lived under the different treatment policies. Treatments are assumed to reduce the probability of weight gain and increase the probability of weight loss.

Additionally, death rates are lower for individuals with lower weights. Consequently, we would expect to have higher effectiveness (more years lived) under the treatment policies, and this expectation is supported by our results. As is usually the case, treatments generate benefits, and also cost more. A standard approach for comparing treatments involves computing the cost savings per unit of effectiveness gained (the ICER). The ICERs (relative to the baseline of *Treat None*) are given in the third column. In the fourth column, we calculate ICERs between successive pairs of options when options are ranked in order of increasing effect (after eliminating options that are dominated or subject to extended dominance). Figure 5 contains three tables, one for each of the three cases considered—no social influence (influence factor of 1), medium social influence (influence factor=1.5) and high social influence (influence factor of 2).

We observe first that there are also cost savings involved in not treating anyone. This is because people who are not treated die younger, and costs terminate with death. Additionally, there are no costs for treatment when people are alive in the no-treatment group. Treatment leads people to live longer, and incur costs for a longer time. Our cost numbers combine the costs from this effect with the costs of the treatment plan. Costs are highest with the *Treat All* plan, under which all obese and overweight individuals are treated. As would be expected, the *Treat Boundary Spanners* plan costs less than *Treat All*. Effectiveness (life expectancy) is highest under *Treat All*, is lowest for *Treat None*, with *Treat Boundary Spanners* in between. We note that, for the two treatment options, effectiveness increases with the influence factor. For *Treat All*, costs decrease somewhat with the influence factor (as a consequence of individuals being in lower weight categories). However, costs of *Treat Boundary Spanners* are not very different. Most of the gain for this treatment plan comes from increased effectiveness (lower mortality). The effect of the strength of social influence on the weight distribution is displayed in figure 6. We observe here that people are in lower weight categories with treatment, and this effect is accentuated when social influence is stronger.

To evaluate cost effectiveness, we first consider the ICER relative to the baseline of no treatment for each of the treatment options (column 3). This is relevant for evaluation when, in addition to the baseline, only one treatment option is feasible (eg, *Treat None* vs *Treat All*). When all three options are feasible, a more detailed incremental analysis is warranted. For this we include the ICER computed for successive options (in column 4). For pairwise comparisons, we need to determine whether the ICERs are less than some acceptable threshold. When all three options are available then, in the no social influence case, *Treat Boundary Spanners* is eliminated since it is subject to extended dominance. What remains is a pairwise comparison and we would need to judge whether \$11 996 is an acceptable increase in cost for the gain of a

Figure 5 Cost effectiveness and incremental cost-effectiveness ratios.

No Social Influence (Influence Factor = 1)				
Treatment Group	Cost	Effectiveness	ICER (pairwise)	ICER
None	\$378,963	56.552		
Boundary Spanners	\$406,240	58.81	\$12,081 / year	*
All	\$435,052	61.227	\$11,996 / year	\$11,996/year

* Eliminated in view of extended dominance of the strategy

Medium Social Influence (Influence Factor = 1.5)				
Treatment Group	Cost	Effectiveness	ICER (pairwise)	ICER
None	\$378,364	56.455		
Boundary Spanners	\$408,485	60.761	\$6,995 / year	\$6995/year
All	\$429,451	62.993	\$7814 / year	\$9393/year

High Social Influence (Influence Factor = 2)				
Treatment Group	Cost	Effectiveness	ICER (pairwise)	ICER
None	\$379,252	56.844		
Boundary Spanners	\$407,745	61.637	\$5945 / year	\$5945/year
All	\$423,654	63.443	\$6728 / year	\$8809/year

year of life. In the medium social influence case, if an acceptable threshold lies between \$6995/year and \$9393/year, then the optimal decision would be *Treat Boundary Spanners*, whereas if the acceptable threshold exceeds \$9393/year, then the optimal decision would be *Treat All*. In the former case, the additional gains in mortality are not worth the incremental cost of treating everyone, whereas in the latter case they are. Similar considerations apply in the high social influence case.

Comparing the ICERs, we find that cost effectiveness increases with the influence factor. In fact, when the influence factor is 2 the ICER (pairwise) for each treatment policy is about half of the value in the no social influence case. This shows that *social influence can have significant effects on the cost effectiveness of treatment policies*. Interestingly, we find that (relative to the no social influence case) the cost-effectiveness rankings become reversed. This is a consequence of the fact that *Treat Boundary Spanners* is subject to extended dominance in the no social influence case, but not when social influences are present. Hence, when the influence factor is 1, *Treat All* is more cost effective than *Treat Boundary Spanners*. However, in the other two cases *Treat Boundary Spanners* is more cost effective—it is preferred at thresholds between \$6995/year and \$9393/year when social influence is medium, and between \$5945/year and \$8809/year when it is high. In other words, for small acceptable thresholds, the option of only treating boundary spanners would be selected over the option of treating everyone. There are values of the acceptable threshold (eg, \$6000/year) for which a treatment policy (*Treat Boundary Spanners*) would be chosen only if social

influences are strong enough (Influence Factor=2). This demonstrates the fact that *optimal treatment policies can be designed to take network structure into account*. Here, in the presence of network effects, we find that focusing treatment only on individuals who occupy key positions in the network is more cost effective than treating everyone. Under stringent standards, the former policy would be acceptable whereas the latter would not be.

Finally, in figures 7 and 8, we examine some effects of variations in the network structure. Figure 7 reports the ICERs for the two treatment policies when the policy is compared to the baseline of no treatment. We only report results for when social influences are present. When social influence is absent changing network structure will have no effect. We vary two aspects of network structure—the number of groups to which an individual can belong (N), and the definition of which boundary spanners to target. The latter is defined via the ‘spanning threshold’, that is, the number of groups to which an individual must belong, before they are targeted for treatment (P). N is a measure of social fragmentation—as it increases, society is less fragmented, with individuals having social association with many different groups. P allows us to identify individuals by the number of others they influence (but it is also true that they are comparatively immune to social influence, since they interact with so many individuals in different groups). Figure 8 reports the ICER computed for successive options after elimination of dominated and extendedly dominated options.

The numbers in the first row of figures 7 and 8 correspond to the baseline case (where the maximum number

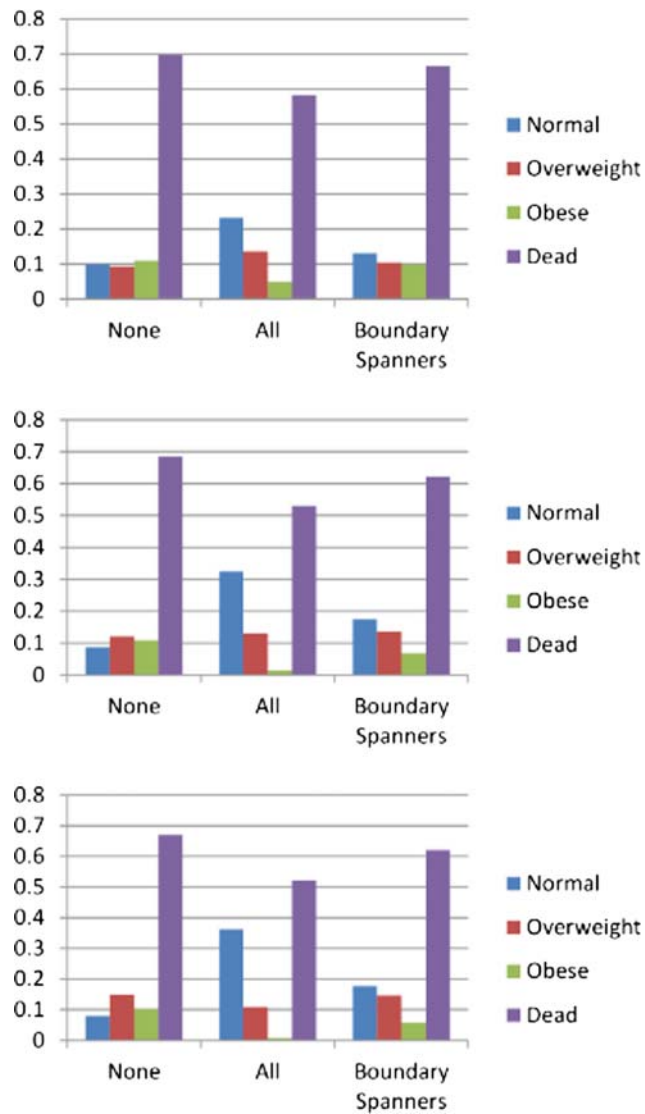


Figure 6 Weight distribution at the end of the simulation when there is no social influence (top), when the social influence is medium (middle) and when social influence is high (bottom).

of groups is $N=3$, and the spanning threshold is $P=2$) reported in figure 5. In the rest of the tables, we allow the number of groups to which an individual could belong to be 3, 5 or 10, and allow the spanning threshold to vary as well. When the number of groups is 3, the only possibilities for the threshold are 2 and 3. When the number of groups is 5, we allow the threshold to take all possible values. When the number of groups is 10, we consider values 2, 3, 4, 5 and 6. For fixed N , higher values of P affect only the *Treat Boundary Spanners* strategy and involve a more restrictive choice of individuals treated. Results are now summarised:

1. Ranking of treatments: When $N=3$, it is most cost effective to treat only the boundary spanners (for $P=3$, at thresholds between \$6599/year and \$8680/year in the medium influence case, and between \$6580/year and \$6827/year in the strong influence

ICER (pairwise)				
	Medium Influence		Strong Influence	
Network (N-P)	Treat Boundary Spanners	Treat All	Treat Boundary Spanners	Treat All
3-2	\$6,995/year	\$7,814/year	\$5,945/year	\$6,728/year
3-3	\$6,593/year	\$7,795/year	\$6,578/year	\$6,713/year
5-2	\$7,812/year	\$7,967/year	\$6,292/year	\$7,433/year
5-3	\$7,653/year	\$7,968/year	\$7,732/year	\$7,437/year
5-4	\$8,252/year	\$7,981/year	\$10,719/year	\$7,452/year
5-5	\$9,111/year	\$7,960/year	\$11,921/year	\$7,429/year
10-2	\$7,791/year	\$8,260/year	\$11,806/year	\$10,217/year
10-3	\$7,827/year	\$8,284/year	\$12,302/year	\$9,904/year
10-4	\$7,901/year	\$8,276/year	\$12,815/year	\$10,478/year
10-5	\$8,433/year	\$8,271/year	\$14,673/year	\$10,605/year
10-6	\$9,469/year	\$8,273/year	\$15,914/year	\$11,049/year

Figure 7 Effects of changing network structure ($N-P$ denotes the structure where the maximum number of groups is N and the spanning threshold is P). The table depicts incremental cost effectiveness ratios relative to the baseline of no treatment.

case). When $N=5$, this strategy is more cost effective for small values of P : for instance, when $P=2$, it is cost effective at thresholds between \$7813/year and \$8293/year with medium influence, and between \$6293/year and \$10,878/year with strong influence. Treating everyone becomes more cost effective for large P (at thresholds above the ICERs in the *Treat All* column in figure 8). The switching point varies with strength of influence. We see this in figure 7 where the pairwise ICER eventually becomes larger in the

ICER				
	Medium Influence		Strong Influence	
Network (N-P)	Treat Boundary Spanners	Treat All	Treat Boundary Spanners	Treat All
3-2	\$6,995/year	\$9,393/year	\$5,945/year	\$8,809/year
3-3	\$6,599/year	\$8,680/year	\$6,580/year	\$6,827/year
5-2	\$7,813/year	\$8,293/year	\$6,293/year	\$10,878/year
5-3	\$7,652/year	\$8,409/year	*	\$7,437/year
5-4	*	\$7,981/year	*	\$7,452/year
5-5	*	\$7,960/year	*	\$7,429/year
10-2	\$7,791/year	\$9,328/year	*	\$10,217/year
10-3	\$7,828/year	\$9,171/year	*	\$9,904/year
10-4	\$7,903/year	\$8,876/year	*	\$10,478/year
10-5	*	\$8,271/year	*	\$10,605/year
10-6	*	\$8,273/year	*	\$11,049/year

Figure 8 Effects of changing network structure ($N-P$ denotes the structure where the maximum number of groups is N and the spanning threshold is P). The table depicts incremental cost-effectiveness ratios. *Denotes choices subject to extended domination.

- Treat Boundary Spanners* case (equivalently, this strategy is eventually subject to extended domination). When $N=10$, this switching pattern continues to hold when influence is medium, but *Treat All* is always more cost effective when influence is strong (at all thresholds above the ICERs in the *Treat All* column of figure 8).
- Increasing the spanning threshold (P): For the *Treat All* policy, changes in P will naturally have no effect on pairwise ICERs (the numbers differ only because of random variations across runs and, even then, are very close). In the *Treat Boundary Spanners* case, all pairwise ICERs increase when social influences are *strong*. Comparing with *Treat All*, treating individuals who meet a spanning threshold of 2 is most cost effective for $N=3$ and 5. So, for instance, at a threshold of \$6000/year only individuals with a spanning threshold of 2 are treated. Treating everyone is most cost effective when $N=10$ (for instance, at thresholds above \$11 000/year). When social influences are *medium*—for *Treat Boundary Spanners*, pairwise ICERs decline for the $N=3$ case, increase in the $N=10$ case, and first decline and then increase when $N=5$. Correspondingly, the most cost-effective strategies are to treat those who meet a spanning threshold of 3 when $N=3$ (at a threshold of \$6600/year for instance) or $N=5$ (eg, at a threshold of \$7700/year), and treat those with a spanning threshold of 2 when $N=10$ (eg, at a threshold of \$7800/year). The reason underlying this pattern is that when P is increased, fewer people meet the treatment threshold. This has the effect of reducing costs, and also effectiveness. The tables depict the net effect on ICERs. This effect depends in complex ways on the nature of the network, the strength of social influences and the parameters of the simulation.
 - Increasing the number of groups (N): For fixed values of P and influence, pairwise ICERs increase when N is larger. The reason is that, for fixed P, when N is large there are many more boundary spanning individuals (so costs are higher). While this tends to increase effectiveness as well, this effect is moderated by the fact that people are not as easily influenced by individuals with a large number of connections (since, on average, they belong to many groups). The net effect is an increase in ICERs.
 - Increasing influence: We hold treatment, N, and P fixed, while changing influence from medium to strong. For *Treat All*, in the case of $N=3$ and 5, the pairwise ICER decreases, whereas it increases for $N=10$. For *Treat Boundary Spanners*, pairwise ICERs decrease for $N=3$, and increase for $N=10$. For $N=5$, pairwise ICERs decrease when P is small, but increase when P is large. This pattern is primarily a consequence of the fact that costs increase with influence when $N=10$, and for large P when $N=5$ (while costs decrease in the remaining cases). This is, in turn, a consequence of the fact that more individuals are obese in such instances (eg, increase in influence when $N=10$), incurring greater cost.

In conclusion, we summarise the results of the sensitivity analysis. The following enumeration corresponds to our list of sensitivity analyses from the Methodology section. (1) Upon reducing (increasing) costs, we found that cost-effectiveness measures (the ICERs) decreased (increased) while the ranking of treatments remained unchanged. (2) We also found that increasing the effectiveness of treatments makes the treatment policies more cost effective. Since both these results are unsurprising, these tables are omitted. (3) We report the results of a sensitivity analysis that involved changing the initial assignment of individuals to groups. In the analysis above, the assignment of individuals to groups is random and does not take their weight into account. This means, in effect, that there is no initial segregation by weight. We relax this assumption and examine a case where there is perfect segregation by weight categories. We then consider the same three policies and report results in figure 9. We find that the *Treat Boundary Spanners* policy becomes much more cost effective in this case when social influences are present. When social influence is medium, the *Treat Boundary Spanners* policy is cost effective for thresholds between \$4813 and \$11 661. When social influence is strong, the policy is cost effective for thresholds between \$3726/year and \$15 947/year. (4) The effects of changing network structure were discussed above and are reported in Tables 7 and 8.

LIMITATIONS

To keep the analysis tractable, and because relevant data are unavailable, we make a number of simplifying assumptions. Here we enumerate the assumptions, and also discuss how these assumptions might be relaxed. On the design of the network, it would be desirable to match this to hard data about the nature of ties in the relevant population. As it is, a lot of studies have focused on the Framingham data, and more information is needed on health-relevant ties. Ideally, for an appropriate sample from the target population, we would map the web of social influences. The maintained assumption in this paper has been that ties are both homogeneous and bi-directional (eg, mutual friendship, and not one-way admiration). It is to be expected that different types of ties (eg, family, friends, co-workers, etc) exert different levels of influence, and future studies should account for this. Network information could most plausibly be collected in the context of studies of youths in the setting of a school, with sociometric surveys being used to keep track of friendship ties, perceptions about role models, relative importance of family versus peers etc. Schools in many communities also have relatively stable populations, so that changes in weight status could be relatively easily recorded and tracked over time. A potential complication is that social ties may evolve over time, but this can be easily accommodated within our framework.

A second limitation here is the model of weight progression. We have used approximations to weight transition probabilities, adjusting only for differences in

Figure 9 Cost effectiveness and incremental cost-effectiveness ratios when individuals are, initially, perfectly segregated into weight categories.

No Social Influence (Influence Factor = 1)				
Treatment Group	Cost	Effectiveness	ICER (pairwise)	ICER
None	\$378,963	56.552		
Boundary Spanners	\$397,291	58.078	\$12,005 / year	*
Treat All	\$435,052	61.227	\$11,996 / year	\$11,996/year

* Eliminated in view of extended dominance of the strategy

Medium Social Influence (Influence Factor = 1.5)				
Treatment Group	Cost	Effectiveness	ICER (pairwise)	ICER
None	\$376,654	55.651		
Boundary Spanners	\$397,563	59.995	\$4813 / year	\$4813/year
All	\$430,355	62.807	\$7504 / year	\$11661/year

High Social Influence (Influence Factor = 2)				
Treatment Group	Cost	Effectiveness	ICER (pairwise)	ICER
None	\$376,512	55.573		
Boundary Spanners	\$397,719	61.264	\$3726 / year	\$3726/year
All	\$427,461	63.129	\$6743 / year	\$15947/year

weights of individuals with whom there are shared ties. It would be desirable to include the effect of variables such as age, sex, race, income, education, etc. A longitudinal study (of the kind described in the previous paragraph) could be used to assess the probability of weight transitions after conditioning for such demographic variables. With such information, treatment can be tailored based not only on positions in networks, but also age, sex, socioeconomic status, etc. Similarly, instead of just three weight classifications, one could construct a model with a larger number of BMI categories (but this would necessitate obtaining the corresponding transition probabilities). One could also disaggregate the health effects of obesity by including additional states (such as diabetic, hypertensive, etc). This would allow us to disaggregate the costs of living with obesity. A related difficulty stems from our Markov assumption, by which transition probabilities depends on the current weight and not weight history. While this assumption can be relaxed within the simulation framework, this poses serious challenges (not least among these is the estimation of transition probabilities).

Third, more refined estimates of social influence would be desirable. The extent of influence could be conditioned on the nature of ties, the frequency of contact and other such variables. A longitudinal study, such as outlined above, can yield meaningful estimates of influence. However, there are significant pitfalls in the use of observational (as opposed to experimental) data to measure social effects. One significant problem of homophily (whereby individuals with similar propensities—in this instance, for weight gain—may choose to associate with one another). Another is the presence of contextual effects whereby individuals within a group

may face common causes influencing weight changes. Christakis and Fowler⁸ suggest ways to control for such problems, but there is disagreement about whether they have effectively done so.²⁴ It is not clear that there is an entirely satisfactory way to overcome the problem short of an experiment which randomises the assignment of individuals into groups. Such an experimental study would allow us to measure the isolated effect of social influence (which is overstated when homophily and contextual effects are not controlled for). Resulting estimates would more accurately reflect the causal effect of others' weight on an individual's propensity to gain weight. Additional studies on the extent of homophily, and common contextual influences (such as targeted marketing to certain groups), would allow us to more fully specify the effects that different treatment policies will have.

Finally, we are aware that the probability of death in any of the three weight categories is not a simple linear equation. However, given our focus on qualitative conclusions, this seems adequate. This is particularly so because we are comparing levels of social influence to our baseline, which was given the same probabilities for death.

CONCLUSION

In recent years, there has been growing realisation of the importance of social influence as a determinant of an individual's health. Separately from this research, there is also now increased interest in the cost effectiveness of treatments. In the USA, the recently passed healthcare reform legislation will only increase the pressure on the system to come up with ways to contain cost. Hence, interest in cost effectiveness is likely to persist, if not increase. The current paper is at the intersection of

these two literatures. We integrate the CEA methodology (specifically, Markov models) with social network analysis to present a tractable approach towards the evaluation of treatments when social influences are present. We show that social influences can have significant effects, and that treatment policies can be devised to take network structure into account (eg, treating only boundary spanning individuals).

There are many avenues for future research. Some of these are detailed in the previous section, where we outline the limitations our obesity prevention study and suggest improvements. There is no reason why the methodology should not extend to problems other than obesity. There are many well-articulated models of disease progression in the literature. It should be possible to combine these with a more detailed specification of networks, and assumptions about how information and behaviours spread through networks. A difficulty is that we often do not have adequately detailed information on individuals' social networks (or the strength of social influence). As discussed earlier, this problem can be overcome if we gather network data as part of the design of studies. However, even when such information is not available, we can use the wealth of available knowledge on *qualitative properties* of networks and how these relate to the spread of behaviours. A more detailed mapping of such qualitative properties to health outcomes and to cost effectiveness would be very helpful for the design of optimal policies. The strategy of using positions in networks (boundary spanning) to target treatment was utilised in this paper. There is a need to explore other networks, and other targeting strategies.

Finally, even when information on social networks is unavailable correlations between observable characteristics of individuals (such as age, occupation or ethnic group) and their position in networks could be useful. For instance, if the young are more likely to be boundary spanners then, even without specific information about individuals' positions in networks, treatment policies can be devised based on age.

Acknowledgements We thank Dr Ritu Agarwal and Mr P Kenyon Crowley for helpful discussions. We are grateful to the reviewers for their thoughtful suggestions.

Contributors Both authors conceived and designed the study. CK programmed and ran the simulations. KP took the lead in drafting the manuscript. Both authors have reviewed the contents and approved the manuscript.

Competing interests None.

Provenance and peer review Not commissioned; externally peer reviewed.

Data sharing statement There are no unpublished data. The computer programs used for the simulations are available, free of charge, from CK.

REFERENCES

1. Berkman LF, Kawachi I, eds. *Social epidemiology*. New York: Oxford University Press, 2000.

2. Levy JA, Pescosolido BA. *Social networks and health*. London: Elsevier, 2002.
3. Smith K, Christakis N. Social networks and health. *Annu Rev Sociol* 2008;34:405–29.
4. Berkman L. Social support, social networks, social cohesion and health. *Social Work Health Care* 2000;31:3–14.
5. Borgatti SP, Mehra A, Brass DJ, et al. Network analysis in the social sciences. *Science* 2009;323:892–5.
6. Epstein J, Longini I, Halloran M, et al. Containing a large bioterrorist smallpox attack: a computer simulation. *Int J Infect Dis* 2007a;11:98–108.
7. Epstein J, Bobashev G, Goedecke D, et al. Controlling pandemic influenza: the role of international travel restrictions. *Pub Library Sci One J* 2007b;2:e401.
8. Christakis N, Fowler J. The spread of obesity in a large social network over 32 years. *N Engl J Med* 2007;357:370–9.
9. Hammond RA. Social influence and obesity. *Curr Opin Endocrinol, Diabetes Obes* 2010;17:467–71.
10. Burke M, Heiland F, Nadler C. From 'overweight' to 'about right': evidence of generational shift in body weight norms. *Obesity* 2010;18:1226–34.
11. Hammond RA. Complex systems modeling for obesity research. *Prev Chronic Dis* 2009;6. http://www.cdc.gov/pcd/issues/2009/jul/09_0017.htm. (accessed Dec 2010)
12. Salvy SJ, Howard M, Read M, et al. The presence of friends increases food intake in youth. *Am J Clin Nutr* 2009;90:282–7.
13. Bahr DB, Browning RC, Wyatt HR, et al. Exploiting social networks to mitigate the obesity epidemic. *Obesity* 2009;17/4:1930–7381.
14. Christakis N, Fowler J. The collective dynamics of smoking in a large social network. *N Engl J Med* 2008;358:2249–58.
15. Gold MR, Siegel JE, Russell LB, et al. *Cost-effectiveness in health and medicine*. New York: Oxford University Press, 1996.
16. Drummond MF, et al. *Methods for the economic evaluation of health care programmes*, 3rd edn. Oxford: Oxford University Press, 2005.
17. Chan PS, Nallamothu BK, Gurm HS, et al. Incremental benefit and cost-effectiveness of high-dose statin therapy in high-risk patients with coronary artery disease. *Circulation* 2007;115:2398–407.
18. Choudhry NK, Patrick AR, Antman EM, et al. Cost-effectiveness of providing full coverage to increase medication adherence in post myocardial infarction medicare beneficiaries. *Circulation* 2008;117:1261–8.
19. Post SE, Sodhi NK, Peng C, et al. A simulation shows that early treatment of chronic hepatitis B infection can cut deaths and be cost-effective. *Health Affairs* 2011;30:340–8.
20. Cawley J. The economics of childhood obesity. *Health Affairs* 2010;29:364–71.
21. Roux L, Tengs TO, Yore MM, et al. Are public health efforts to promote physical activity cost-effective? A cost-effectiveness analysis of the Guide to Community Preventive Services recommendations for increasing physical activity. *Am J Prev Med* 2008;35:578–88.
22. Valente TW. *Social networks and health: models, methods, and applications*. Oxford: Oxford University Press, 2010.
- 22a. Borgatti SP, Everett MG, Freeman LC. UCINET for Windows: Software for Social Network Analysis. Analytic Technologies, Harvard, 2002.
23. Hill AL, Rand DG, Nowak MA, et al. Infectious disease modeling of social contagion in networks. *PLoS Comput Biol* 2010;6:e1000968.
24. Cohen-Cole E, Fletcher J. Is obesity contagious? Social networks vs. environmental factors in the obesity epidemic. *J Health Econ* 2008;27:1382–7.
25. Finkelstein EA, Trogon JG, Cohen JW, et al. CDC weight of the Nation Press Briefing. CDC Online Newsroom. July 27, 2009, 11:30 a.m. ET. <http://www.cdc.gov/media/transcripts/2009/t090727.htm> (accessed 29 October 2009).
26. Finkelstein EA, Trogon JG, Cohen JW, et al. Annual Medical Spending Attributable To Obesity: Payer-And Service-Specific Estimates. *Health Affairs* 2009;28:w822–w831.
27. Hammond RA, Levine R. The economic impact of obesity in the United States. *Diabetes Metab Syndr Obes* 2010;3:285–95.
28. Maheswaran H, Petrou S, Rees K, et al. Estimating EQ-5D utility values for major health behavioural risk factors in England. *J Epidemiol Community Health* 2012. In press. doi:10.1136/jech-2012-201019.
29. Hoffmann Rose L. Costly calories. *Forbes Magazine*, 6 April. 2005, http://www.forbes.com/2005/04/06/cx_lrlh_0406costlycalories.html (accessed 29 October 2009).