Econometric analysis of models with social interactions

Brendan Kline (UT-Austin, Econ) and Elie Tamer (Harvard, Econ)

2017 - Please see the paper for more details, citations, etc.

The setting

- Social interactions or peer effects, or spillovers, or ...: My outcome/choice affects my friends' outcomes/choices, and vice versa
- Applications:
  - alcohol use (e.g., Kremer and Levy (2008)), smoking (e.g., Powell, Tauras, and Ross (2005)), other substance use (e.g., Lundborg (2006))
  - education outcomes (e.g., Epple and Romano (2011), Sacerdote (2011))
  - obesity (e.g., Christakis and Fowler (2007))
- Similar models can apply to other settings with interacting decision makers:
  - Firm behavior United Airlines v. American Airlines...
  - etc.

The questions

- What is the statistical model?
  - Linear-in-means model: a linear model with an extra "spillover" term
  - Non-linear models: models based on game theory?
- How is the model identified? Under what conditions can the parameters of the model be estimated?
  - Linear-in-means model: the reflection problem (e.g., Manski (1993))
  - Non-linear models: multiple equilibria Our model may not predict a unique outcome. Example: among friends, it is an equilibrium if we all smoke, or none of us smoke.
- What should we do with the estimates? What variables do we think policy can manipulate?

Social networks

- Many (most?) models of social interactions are built on social network data (an interesting topic by itself)

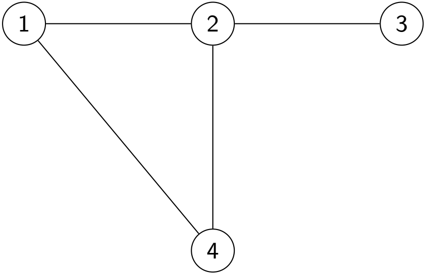

Social networks

- Equivalently, as the *(symmetric) adjacency matrix*:

$$g = \begin{pmatrix} 0 & 1 & 0 & 1 \\ & 0 & 1 & 1 \\ & & 0 & 0 \\ & & & 0 \end{pmatrix}$$

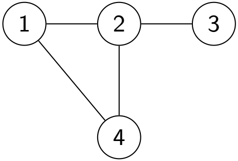

  where g_{ij} element in row i and column j indicates whether i and j are linked in the network

- We are going to assume this is in the dataset. Example: AddHealth dataset surveys high school students. Other examples: Facebook friends. Twitter interactions. etc.

- Sometimes we only know group membership, but not friendships within a group. In that case, we *might* be willing to set g to show links among all individuals.

The linear-in-means model

- Some variables:
  - y_{ig} is the outcome/decision of individual i in group g (e.g., # of cigarettes, test score)
  - x_{ig} are observed characteristics of individual i in group g (e.g., demographics, parental characteristics)
  - ϵ_{ig} are unobserved characteristics of individual i in group g
- Collect data on many groups, and the individuals within each group (e.g., students in different schools)
- The linear-in-means model (e.g., Manski (1993)):

$$y_{ig} = \alpha + x_{ig}\beta + \left(\sum_{j\neq i} g_{ij}^{w} x_{jg}\right)\gamma + \left(\sum_{j\neq i} g_{ij}^{w} y_{jg}\right)\delta + \epsilon_{ig}$$

  - Delete the red and blue terms $\Rightarrow$ an ordinary linear model
  - Include the red and blue terms $\Rightarrow$ allow y_{ig} to depend on the x's and y's of other people in the same group

Linear-in-means model

- How to interpret $\left(\sum_{j\neq i} g_{ij}^{w} x_{jg}\right)\gamma + \left(\sum_{j\neq i} g_{ij}^{w} y_{jg}\right)\delta$...
- What is g_{ij}^{w}?
  - Commonly but not always,

$$g_{ij}^{w} = \begin{cases} 0 & \text{if } g_{ij} = 0 \\ \frac{1}{\sum_{k\neq i} g_{ik}} & \text{if } g_{ij} = 1 \end{cases}.$$

- Therefore, $\left(\sum_{j\neq i} g_{ij}^{w} x_{jg}\right)$ is the *weighted average* of the x's of the other individuals in the group
- The "weights" are equal for all friends, and 0 for non-friends
- And $\left(\sum_{j\neq i} g_{ij}^{w} y_{jg}\right)$ is the *weighted average* of the y's of the other individuals in the group
- So the model for y_{ig} is **linear** (obviously) **in the means** of these characteristics of the others in the group

Identification in the linear-in-means model

- How can we think about estimating the parameters of

$$y_{ig} = \alpha + x_{ig}\beta + \left(\sum_{j \neq i} g_{ij}^{w} x_{jg}\right)\gamma + \left(\sum_{j \neq i} g_{ij}^{w} y_{jg}\right)\delta + \epsilon_{ig}?$$

- We **cannot** just "run this regression" since the outcomes are simultaneously determined ... y_{ig} affects y_{jg} *(potentially... depending on g)* so there is "reverse causality" in this equation for y_{ig}
- We need to write the model at the group level:
  - Y_g is a vector that stacks the elements y_{ig}
  - X_g is a matrix that stacks the vectors x_{ig} for different individuals in different rows
  - ϵ_g is a vector that stacks the elements ϵ_{ig}
  - G^w is a matrix of weighted social influence, with g_{ij}^w in row i and column j
- Then:

$$Y_g = 1_{N_g \times 1}\alpha + X_g\beta + G^w X_g \gamma + G^w Y_g \delta + \epsilon_g.$$

Identification in the linear-in-means model

- Solve

$$Y_g = 1_{N_g \times 1}\alpha + X_g\beta + G^w X_g \gamma + G^w Y_g \delta + \epsilon_g$$

for the "reduced form" that puts the endogenous variables (the Y's) on the LHS, and everything else on the RHS:

$$Y_g = (I - \delta G^w)^{-1}\left(1_{N_g \times 1}\alpha + X_g\beta + G^w X_g \gamma + \epsilon_g\right).$$

- The inverse exists as long as $-1 < \delta < 1$. Generally, $\delta > 0$, so $\delta < 1$ means increasing the average outcome of friends by 1 increases your outcome by *less than* 1.

Identification in the linear-in-means model

- Using the exogeneity assumption $E(\epsilon_g|X_g, G^w) = 0$, it follows that

$$E(Y_g|X_g, G^w) = (I - \delta G^w)^{-1}\left(1_{N_g\times 1}\alpha + X_g\beta + G^w X_g\gamma\right),$$

  and therefore

$$E(y_{ig}|X_g, G^w) = e_i'(I - \delta G^w)^{-1}\left(1_{N_g\times 1}\alpha + X_g\beta + G^w X_g\gamma\right),$$

  where e_i is the unit vector, with 1 as the i-th element and 0 as every other element.
- The data reveals the LHS directly, so the question is whether it is possible to “solve” (in some sense...) the equation to learn the parameters on the RHS.

Identification in the linear-in-means model

- Useful to think about implications marginal effects of the model

$$y_{ig} = \alpha + x_{ig}\beta + \left(\sum_{j \neq i} g_{ij}^{w} x_{jg}\right)\gamma + \left(\sum_{j \neq i} g_{ij}^{w} y_{jg}\right)\delta + \epsilon_{ig}$$

- Because the parameters relate to the "effects", it makes sense to study the marginal effects of the components of X_g:

$$\frac{\partial E(y_{ig} | X_g, G^w)}{\partial X_g} = \beta e_i'(I - \delta G^w)^{-1} + \gamma e_i'(I - \delta G^w)^{-1} G^w.$$

- If $\gamma = 0$ and $\delta = 0$, then the marginal effect is simply $\frac{\partial E(y_{ig}|X_g, G^w)}{\partial X_g} = \beta e_i'$, which is a complicated way to say the model behaves like an ordinary linear model (obvious)

Identification in the linear-in-means model

- Useful to think about implications marginal effects of the model

$$y_{ig} = \alpha + x_{ig}\beta + \left(\sum_{j \neq i} g^w_{ij} x_{jg}\right)\gamma + \left(\sum_{j \neq i} g^w_{ij} y_{jg}\right)\delta + \epsilon_{ig}$$

- Because the parameters relate to the "effects", it makes sense to study the marginal effects of the components of X_g:

$$\frac{\partial E(y_{ig}|X_g, G^w)}{\partial X_g} = \beta e_i'(I - \delta G^w)^{-1} + \gamma e_i'(I - \delta G^w)^{-1} G^w.$$

- If $\delta = 0$, then the marginal effect is $\frac{\partial E(y_{ig}|X_g, G^w)}{\partial X_g} = \beta e_i' + \gamma e_i' G^w$, which is a complicated way to say the x's have two effects:
  - My x affects my y, per β
  - My friends x's affect my y, per γ

Identification in the linear-in-means model

- Useful to think about implications marginal effects of the model

$$y_{ig} = \alpha + x_{ig}\beta + \left(\sum_{j \neq i} g_{ij}^{w} x_{jg}\right)\gamma + \left(\sum_{j \neq i} g_{ij}^{w} y_{jg}\right)\delta + \epsilon_{ig}$$

- Because the parameters relate to the "effects", it makes sense to study the marginal effects of the components of X_g:

$$\frac{\partial E(y_{ig}|X_g, G^w)}{\partial X_g} = \beta e_i'(I - \delta G^w)^{-1} + \gamma e_i'(I - \delta G^w)^{-1} G^w.$$

- In general, the x's have three interrelated effects:
  - My x affects my y, per β
  - My friends x's affect my y, per γ
  - My x affects my y and my y affects my friends' y, per δ
    - Symmetrically, my friends' x's affect my friends' y's, which affect their friends' y's, etc. etc.
    - This is sometimes called the *social multiplier effect* and is an important reason to study social interactions models

Identification in the linear-in-means model

- The marginal effects of the components of X_g:

$$\frac{\partial E(y_{ig}|X_g, G^w)}{\partial X_g} = \beta e_i'(I - \delta G^w)^{-1} + \gamma e_i'(I - \delta G^w)^{-1} G^w.$$

- Suppose that all individuals within a group are linked, maybe because you cannot do better given data limitations
- Then G^w has a particularly simple (highly symmetric) form, namely, it is $\frac{1}{N-1}$ in all non-diagonal entries, which lets you prove that:
  - The effect of the k-th explanatory variable of individual i on the outcome of individual i is the same as the effect of the k-th explanatory variable of individual j on the outcome of individual j, for all k and i, j.
  - The effect of the k-th explanatory variable of individual j on the outcome of individual i is the same as the effect of the k-th explanatory variable of individual m on the outcome of individual l, for all k and $i \neq j$ and $l \neq m$.
- So there are essentially only $2K$ distinct effects of the explanatory variables, but $2K + 1$ parameters $\Rightarrow$ cannot identify (estimate) the parameters, because there is "too much collinearity" of marginal effects

Identification in the linear-in-means model

- Another way to see the problem when $g^w_{ij} = \frac{1}{N-1}$:

$$y_{ig} = \alpha + x_{ig}\beta + \left(\sum_{j \neq i} \frac{1}{N-1} x_{jg}\right)\gamma + \left(\sum_{j \neq i} \frac{1}{N-1} y_{jg}\right)\delta + \epsilon_{ig}$$

  is a system of simultaneous equations with no valid instrumental variables since there is no component of X that affects only y_{ig}.
  - Put differently, x_{ig} affects y_{ig} because of β but also all of the other y_{jg}'s because of γ... so it does not satisfy the exclusion restriction needed for an IV.
- So we need to find conditions on the social network that results in the model having valid *exclusion restrictions* (or *instrumental variables*)

Identification in the linear-in-means model

- It seems the most popular approach concerns *intransitive triads* in the social network (e.g., Bramoullé, Djebbari, and Fortin (2009), De Giorgi, Pellizzari, and Redaelli (2010)).
- Intransitive triad:

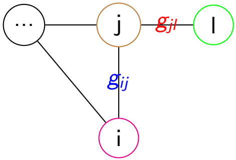

- So, then:
  - x_{lg} does not appear directly in the structural form of the outcome y_{ig},
  - but x_{lg} are relevant instruments for the outcome y_{jg} in the structural form of the outcome y_{ig},
    - since x_{lg} directly affects y_{lg}, which directly affects y_{jg} since $g_{jl} = 1$, which directly affects y_{ig} since $g_{ij} = 1$.
- *Existence of intransitive triads implies valid instruments, so the model parameters are identified* ... requires collecting the social network data

When do we believe the linear-in-means model?

- "All models are wrong, but some are useful" (approximations) ... so in what cases is the linear-in-means model a useful approximation?
- Suppose that the individuals are *playing a game* in which their utility functions are

$$u_{ig}(y_{ig}, y_{-ig}) = \theta_{ig} y_{ig} - \frac{(1-\delta) y_{ig}^2}{2} - \underbrace{\frac{\delta}{2}(y_{ig} - \zeta_{ig})^2}_{\text{unhappy if deviating from average}},$$

  with $\theta_{ig} = \alpha + x_{ig}\beta + \left(\sum_{j \neq i} g_{ij}^w x_{jg}\right)\gamma + \epsilon_{ig}$ and $\zeta_{ig} = \sum_{j \neq i} g_{ij}^w y_{jg}$
- By taking first order conditions, the linear-in-means model describes the *best response functions*
- $\Rightarrow$ If people try to "emulate" the average outcome/choice of their friends, *and can choose any real number as their outcome/choice*, then the linear-in-means model is sensible

When do we believe the linear-in-means model?

- But what about discrete outcomes/choices. Specifically, binary y?
- In an ordinary linear model

$$y_i = x_i\beta + \epsilon_i,$$

  we might not be too concerned about binary y (the linear probability model)
- How would social interactions with binary y work? Example: smoking among pairs of individuals. One possible utility function says:
  - I don't want to smoke alone. My friend doesn't want to smoke alone.
  - But smoking with friends is fun. So I want to smoke if my friend smokes. And my friend wants to smoke if I smoke.
  - What will actually happen? **We cannot say uniquely!**
    - Both smoke. If my friend smokes, I want to smoke. And vice versa.
    - Neither smoke. If my friend doesn't smoke, I don't want to smoke. And vice versa.
    - So we have *multiple equilibria* of this interaction.
- The linear-in-means model cannot possibly capture this aspect.

Social interactions models for binary outcome/choice

- We can model this interaction as a game:

	don't smoke	smoke
don't smoke	$\begin{pmatrix} 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} 0 \\ x_2\beta + \epsilon_2 \end{pmatrix}$
smoke	$\begin{pmatrix} x_1\beta + \epsilon_1 \\ 0 \end{pmatrix}$	$\begin{pmatrix} x_1\beta + \Delta + \epsilon_1 \\ x_2\beta + \Delta + \epsilon_2 \end{pmatrix}$

- As before, the x's are characteristics of the individuals. And the Δ is the "spillover" effect *on utility*
- Depending on the specification of these quantities, can have *multiple equilibria* $\Rightarrow$ different identification strategy needed
- One main idea: If x's vary sufficiently, the multiple equilibria goes away. i.e., if x_2 is huge, player 2 definitely will smoke, so then player 1 basically has a "standard" binary choice problem. Used in Tamer (2003), Bajari, Hong, and Ryan (2010), Kline (2015), etc. Incorporating social network?

What should we do with the estimates?

- What policy questions can we answer with

$$y_{ig} = \alpha + x_{ig}\beta + \left(\sum_{j\neq i} g^w_{ij} x_{jg}\right)\gamma + \left(\sum_{j\neq i} g^w_{ij} y_{jg}\right)\delta + \epsilon_{ig}?$$

- Sometimes, people focus on the "peer effect" or "spillover effect" parameter by itself.
  - It is not necessarily clear this answers a policy intervention question. Can we directly manipulate the outcomes/choices of person i's friend?
    - In settings other than friendship (say, airlines), government able to directly control outcomes
    - Might think about assigning to new peer group... but Carrell, Sacerdote, and West (2013)
- Might think about manipulating x's, and now spillover effects matter to that effect we saw the formula before

$$\frac{\partial E(y_{ig}|X_g, G^w)}{\partial X_g} = \beta e_i'(I - \delta G^w)^{-1} + \gamma e_i'(I - \delta G^w)^{-1} G^w.$$

- So, the parameters help us answer questions, but we should be careful to use them correctly