Economics Letters 137 (2015) 123–126

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Difference-in-differences techniques for spatial data: Local autocorrelation and spatial interaction☆

Michael S. Delgado [a,*], Raymond J.G.M. Florax [a,b]

[a] Department of Agricultural Economics, Purdue University, 403 W. State Street, West Lafayette, IN 47907, USA

[b] Department of Spatial Economics, Vrije Universiteit Amsterdam, and Tinbergen Institute, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands

HIGHLIGHTS

- Develop a difference-in-differences estimator for spatial data.
- Allow for local spatial interaction in potential outcomes.
- Identify direct and indirect treatment effects.
- Monte Carlo simulations show good finite sample performance.

ARTICLE INFO

Article history:
Received 24 May 2015
Received in revised form
24 September 2015
Accepted 26 October 2015
Available online 30 October 2015

JEL classification:
C21
C53

Keywords:
Difference-in-differences
Monte Carlo simulation
Program evaluation
Spatial autocorrelation
Spatial interaction

ABSTRACT

We consider treatment effect estimation via a difference-in-difference approach for spatial data with local spatial interaction such that the potential outcome of observed units depends on their own treatment as well as on the treatment status of proximate neighbors. We show that under standard assumptions (common trend and ignorability) a straightforward spatially explicit version of the benchmark difference-in-differences regression is capable of identifying both direct and indirect treatment effects. We demonstrate the finite sample performance of our spatial estimator via Monte Carlo simulations.

1. Introduction

The linear difference-in-differences (DID) model is a benchmark tool in the program evaluation literature (e.g., Ashenfelter, 1978; Ashenfelter and Card, 1985). At its core, a treatment effect is the difference between two potential outcomes, with potential outcomes being a function of treatment status (Rubin, 1974). The fundamental problem is that units are never observed in both treated and untreated states (Holland, 1986), and identification requires comparison of treated units to untreated (control) units. In the standard DID design, observations $i = (1, 2, \ldots, n)$ are observed in two time periods, $T \in \{0, 1\}$, and are grouped via $D \in \{0, 1\}$ such that $D_i = 1$ indicates treatment. Given a vector of time-varying covariates, X_{it}, the standard DID equation is:

$$y_{it} = \alpha_0 + \alpha_1 X_{it} + \alpha_2 D_{it} + \alpha_3 T_{it} + \alpha_4 D_{it} T_{it} + \varepsilon_{it}, \quad (1)$$

in which ε_{it} is a mean-zero error term that is uncorrelated with D_{it} and T_{it}. It is straightforward to accommodate additional time periods. The identifying assumptions require correct linear specification of the conditional mean, a homogeneous effect of treatment, and the parallel-trends assumption denoting that absent treatment both treated and untreated units evolve along the same temporal path. Strong or weak ignorability (unconfoundedness) is assumed as well, implying that treatment assignment is independent of the

☆ This research was supported in part by computational resources provided by Information Technology at Purdue (ITAP)–Rosen Center for Advanced Computing, Purdue University, West Lafayette, IN, USA. The authors would like to thank Arnab Bhattacharjee, Julie Le Gallo, and an anonymous reviewer for excellent comments.

* Corresponding author.
E-mail addresses: delgado2@purdue.edu (M.S. Delgado), rflorax@purdue.edu (R.J.G.M. Florax).

http://dx.doi.org/10.1016/j.econlet.2015.10.035
0165-1765/

outcome y_{it}, eventually conditional on X_{it}. Maintaining these assumptions and suppressing the subscript, the conditional average treatment effect, $ATE(x)$ is:

$$\begin{aligned} ATE(x) = & \{\mathrm{E}[y|X=x, D=1, T=1] \\ & - \mathrm{E}[y|X=x, D=1, T=0]\} \\ & - \{\mathrm{E}[y|X=x, D=0, T=1] \\ & - \mathrm{E}[y|X=x, D=0, T=0]\}, \end{aligned} \tag{2}$$

or the difference in the conditional differences over time between the treated and control units. From Eqs. (1) and (2), we can see that $ATE = \alpha_4$. Typically, Eq. (1) is estimated with ordinary least squares (OLS), and a rejection of $H_0 : \alpha_4 = 0$ via a t-test is evidence of a significant causal effect.

Another key assumption needed for identification is the stable unit treatment value assumption (SUTVA) (Rubin, 1978, 1990) which "implies that potential outcomes for person i are unrelated to the treatment status of other individuals" (Angrist et al., 1996, p. 446). Violation of the SUTVA assumption, referred to as "interference" or "social interaction", invalidates identification of causal effects in the traditional DID setup, and adjustments are clearly needed. It no longer suffices to only consider one's own treatment status, but the treatment status of other units has to be taken into account as well. For example, imagine a program designed to improve crop yields by enhancing farming ability through education or experience. Communication among neighboring farmers may evoke indirect treatment effects, in which treated and untreated farmers benefit from treatment (knowledge) of others. Similar examples are prevalent throughout economics.

We develop a spatial DID model that accounts for the possibility of spatial correlation in treatments and spatial interaction in treatment responses.[1] The term "spatial interaction" is deliberate and refers to the fact that both direct and indirect treatment effects exist, because potential outcomes are not independent. Indirect treatment effects are sometimes referred to as "spillovers" or "network effects", resulting from contagion, displacement, communication, social comparison, or signaling (Gerber and Green, 2012). Such spillovers create spatial correlation in the treatment response. It is important to note that we are explicitly interested in spatial correlation caused by dependence in potential outcomes, as opposed to spatial correlation stemming from spatially correlated treatments or any other sort of spatial sorting process. Spatial interaction in treatment responses requires a plausible causal mechanism to be at work; existing spatial patterns due to spatial sorting need to be excluded as a source of spatially correlated treatment responses. This is akin to the well known ignorability assumption evoked in the traditional DID setup (see Imbens and Wooldridge, 2009).

We assume treatment is binary, and restrict our focus to spatial interaction in the treatment results that are "local", as opposed to "global". This means that the spatial effects are restricted to immediate neighbors, defined on the basis of contiguity or distance.[2] In effect, this amounts to assuming that SUTVA holds outside the bounds of immediate neighbors. In the context of the farming example one might assume that SUTVA holds *between* villages, but local spatial interaction and hence SUTVA violation is allowed for *within* villages. Generally, however, the only requirement for our approach to be valid is that SUTVA holds across some particular dimension of the sample space; it is not necessarily restricted to a discernable within–between situation. Under these general circumstances, a simple extension of the standard DID setup using spatial econometric tools allows us to estimate the *ATE*, which we decompose into a direct ("own") and indirect ("neighbor") treatment effect. Our spatial DID model is straightforward to implement, and simple Monte Carlo exercises illustrate good finite sample performance.

2. Spatial difference-in-differences

Imagine a setup with n cross-sectional, spatial observations observed over two time periods that can represent points or areas. We operationalize local spatial interaction in treatment responses using the spatial lag operator L^s, defined as W_s, where W_s is a $(2n \times 2n)$ block-diagonal row-standardized spatial weights matrix containing non-zero elements for spatial units belonging to contiguity class s (Anselin, 1988). The contiguity class can vary, as long as the number of neighbors is restricted and does not extend to the entire spatial system. Spatial interaction in outcomes is defined as $D_L T = (I + \rho L^s)D \circ T = (I + \rho W_s)D \circ T$, where I is the identity matrix, ρ a spatial autoregressive parameter, and $\circ$ signals element-by-element multiplication (Hadamard product).

One question is whether we need to distinguish random treatment from spatially correlated treatment.[3] In the case of spatial correlation in treatment the SUTVA assumption is not violated. Moreover, even with spatially correlated exogenous covariates the unbiasedness, consistency, and efficiency properties of OLS are unaltered (Anselin, 1988). Hence the standard DID approach leads to proper identification and estimation of the *ATE*, and we do not require spatially explicit reparameterization of Eq. (1).

The situations with spatial interaction in the responses are different because SUTVA is violated and indirect effects should therefore be explicitly modeled. One common approach is to meticulously identify all treatment and control groups and apply a difference-in-difference-in-differences technique (Imbens and Wooldridge, 2009). There are several disadvantages to such an approach. First, the number of groups becomes unwieldy and the estimator becomes inefficient in small samples.[4] Second, since there are no restrictions on the estimated parameters one may be confronted with illogical "bouncing beta's", such as the case where the treatment response of having two treated neighbors may be greater than that of having four treated neighbors. We therefore suggest a spatially structured approach.

Without loss of generality the model with spatial interaction in the responses is:

$$\begin{aligned} y &= \alpha_0 + \alpha_1 X + \alpha_2 D + \alpha_3 T + \alpha_4 D_L \circ T + \varepsilon, \\ &= \alpha_0 + \alpha_1 X + \alpha_2 D + \alpha_3 T + \alpha_4 (I + \rho W) D \circ T + \varepsilon, \\ &= \alpha_0 + \alpha_1 X + \alpha_2 D + \alpha_3 T + \alpha_4 D \circ T + \alpha_5 WD \circ T + \varepsilon, \end{aligned} \tag{3}$$

in which $\alpha_5 = \rho\alpha_4$, and D represents either random or spatially correlated treatments. We assume $\rho \neq 0$ so that the model does not revert to the standard DID equation. Omitting $WD \circ T$ renders

[1] In spatial econometrics, "autocorrelation" and "dependence" are often treated as synonymous, even though dependence is a characteristic of the joint probability density function and can only be measured as correlation under assumptions such as normality, stationarity, etc. Spatial patterns on a map can be generated through a dependence mechanism that may or may not be known, or result from spatial heterogeneity (Anselin, 1988).

[2] In the global case, indirect effects propel through the entire spatial system; everybody is a neighbor of everybody else, with interactions (or correlations) subject to distance decay.

[3] Given that D is binary, join count statistics can be used to determine whether there is spatial autocorrelation in treatment (Cliff and Ord, 1981).

[4] A spatial setting with local indirect effects implies that groups would have to be defined such that units with different numbers of treated neighbors would fall into different treatment or control groups. In the simple case of a (10×10) spatial grid with adjacency defined as the sharing of a border or vertex, the number of groups based on spatially interactive responses is 9×2, the number of treated neighbors $(0, 1, \ldots, 8)$ multiplied by 2 to incorporate whether own treatment is 0 or 1.