

The Spatial Probit Model of Interdependent Binary Outcomes: Estimation, Interpretation, and Presentation

Robert J. Franzese, Jr. (franzese@umich.edu)
Professor of Political Science, University of Michigan, Ann Arbor

Jude C. Hays (jchays@uiuc.edu)
Assistant Professor of Political Science, University of Illinois, Urbana-Champaign

4 March 2009

[Presented at the Annual Meeting of the Public Choice Society, 6 March 2009]

ABSTRACT: Interdependence—i.e., that the outcomes in or actions or choices of some units depend on those in/of others—is substantively and theoretically ubiquitous in and central to binary outcomes of interest across the social sciences. Most empirical applications omit interdependence, however; even theoretical and substantive discussion usually ignores it. Moreover, in the few contexts where spatial interdependence has been acknowledged or emphasized, such as in the social-network and policy-diffusion literatures, models do not fully reflect simultaneity of the outcomes across units and/or the endogeneity of the spatial lags (appropriately) employed to model the interdependence has gone unrecognized. This paper notes and explains some of the severe challenges posed by spatial interdependence in binary-outcome models and then follows recent spatial-econometric advances to suggest two simulation approaches for surmounting the analytically intractable and computationally intense estimation demands of these models, frequentist recursive-importance-sampling (RIS) or Bayesian MCMC. In brief, the complications arise because the endogenous spatial-lag implies the conditional independence that typically yields likelihoods for maximization that simply multiply N univariate distributions will not obtain. With interdependent observations, the likelihood is instead one N -variate joint distribution, and the one N -dimensional cumulative-normal in spatial probit is tremendously more intense to compute than the N cumulative standard-normal distributions of the common probit. After discussing Monte-Carlo comparisons of the performance of these alternative estimators, including our own, we show how to apply the same estimation-by-simulation methods to calculate estimated spatial effects of counterfactual shocks in terms of outcomes or probabilities of outcomes (with associated confidence/credibility regions) rather than in terms of parameter estimates or latent variables only as in prior spatial-probit applications.

The Spatial Probit Model of Interdependent Binary Outcomes: Estimation, Interpretation, and Presentation

ABSTRACT: Interdependence—i.e., that the outcomes in or actions or choices of some units depend on those in/of others—is substantively and theoretically ubiquitous in and central to binary outcomes of interest across the social sciences. Most empirical applications omit interdependence, however; even theoretical and substantive discussion usually ignores it. Moreover, in the few contexts where spatial interdependence has been acknowledged or emphasized, such as in the social-network and policy-diffusion literatures, models do not fully reflect simultaneity of the outcomes across units and/or the endogeneity of the spatial lags (appropriately) employed to model the interdependence has gone unrecognized. This paper notes and explains some of the severe challenges posed by spatial interdependence in binary-outcome models and then follows recent spatial-econometric advances to suggest two simulation approaches for surmounting the analytically intractable and computationally intense estimation demands of these models, frequentist recursive-importance-sampling (RIS) or Bayesian MCMC. In brief, the complications arise because the endogenous spatial-lag implies the conditional independence that typically yields likelihoods for maximization that simply multiply N univariate distributions will not obtain. With interdependent observations, the likelihood is instead one N -variate joint distribution, and the one N -dimensional cumulative-normal in spatial probit is tremendously more intense to compute than the N cumulative standard-normal distributions of the common probit. After discussing Monte-Carlo comparisons of the performance of these alternative estimators, including our own, we show how to apply the same estimation-by-simulation methods to calculate estimated spatial effects of counterfactual shocks in terms of outcomes or probabilities of outcomes (with associated confidence/credibility regions) rather than in terms of parameter estimates or latent variables only as in prior spatial-probit applications.

I. Introduction to Spatial Probit

Many phenomena that social scientists study are inherently, or by measurement, discrete choices. Canonical political-science examples include citizens' vote and turnout choices, legislators' votes, governments' policy-enactments, wars among or within nations, and regime type or transition. In all these political contexts, and widely across the social sciences,¹ substantively and theoretically, the choices/outcomes of/in some units depend on those of other units. Whether and for whom citizens vote depends on whether and how their neighbors or social networks vote; legislators' votes depend on how they expect or observe others to vote; governments' policy choices depend on others' policies *via* competition or learning; nations' internal wars may arise in some part through contagion from others' conflicts; whether and which others join conflicts heavily condition states' entry to and involvement in external wars, international organization and treaties; and regime change at home is often spurred by example, fomentation, or otherwise from abroad.

¹ For an extensive, topically organized bibliography of interdependence studies across the social sciences, see appendix to Franzese & Hays (2009a): <http://www.umich.edu/~franzese/FranzeseHays.Interdependence.IPSA.References.pdf>.

Indeed, interdependence seems almost inherent to *social-science discrete-choices*. Nevertheless, beyond a few topical areas, interdependence in discrete outcomes receives very little theoretical or empirical attention. Perhaps the most-extensive and longest-standing exception in political science surrounds the diffusion of policies or institutions across national or sub-national governments. The study of policy diffusion across U.S. States in particular has deep roots and much contemporary interest.² Similar innovation-learning mechanisms underlie some comparative studies of policy diffusion (Schneider & Ingram 1988; Rose 1993; Meseguer 2004, 2005; Gilardi 2005). Interest in institutional or even regime diffusion, too, is long-standing and much invigorated recently in comparative and international politics. Dahl's (1971) classic *Polyarchy*, e.g., (implicitly) references international diffusion among his list of democracy's eight causes; Starr's "Democratic Dominoes" (1991) and Huntington's *Third Wave* (1991) accord it a central role; and O'Loughlin et al. (1998) and Gleditsch & Ward (2006, 2007) have recently estimated its empirical extent. Eising (2002), Brune et al. (2004), Simmons & Elkins (2004), Brooks (2005), Elkins et al. (2006), Simmons et al. (2006), and others likewise stress international diffusion in recent economic liberalizations.³

The other major area of extensive interest in interdependence is micro-behavioral, where some of the long-standing and recently surging interest in *contextual effects* surrounds effects on respondent behaviors or opinions of aggregates of others'—e.g., those of her region, community, or social network. Within the large *contextual-effects* literature in political behavior (Huckfeldt & Sprague 1993 review), recent work stressing interdependence include Braybeck & Huckfeldt (2002ab), Cho (2003), Huckfeldt et al. (2005), Lin et al (2006), Cho & Gimpel (2007), and Cho & Rudolph (2007).

The substantive range of important spatial-interdependence effects on discrete outcomes extends

² E.g., Crain 1966; Walker 1969, 1973; Gray 1973; Knoke 1982; Caldiera 1985; Lutz 1987; Berry & Berry 1990, 1999; Case et al. 1993; Berry 1994; Rogers 1995; Mintrom 1997ab; Mintrom & Vergari 1998; Mossberger 1999; Godwin & Schroedel 2000; Balla 2001; Mooney 2001; Bailey & Rom 2004; Boehmke & Witmer 2004; Daley & Garand 2004; Grossback et al. 2004; Shipan & Volden 2006; Volden 2006. See also the extended bibliography linked in note 1.

³ For an excellent recent review of these diffusion literatures across political science, see Graham et al. (2008).

well beyond inter-governmental/interstate diffusion and social-network effects, however, spanning the subfields and substance of political science. Inside democratic legislatures, representatives' votes depend on others' (expected) votes. In electoral studies, candidate qualities or strategies and citizens votes and election outcomes in some contests depend on (expectations of) those in others. Outside legislative and electoral arenas, the probabilities and outcomes of coups, revolutions, and/or riots in one unit depend in substantively crucial ways on (expectations of) those in others. In international relations, the interdependence of states' actions essentially defines the subfield. Whether states enter wars, alliances, treaties, or international organizations, e.g., depends greatly on how many and who else (are expected to) enter. Interdependence is substantively crucial in comparative and international political economy too; globalization, for instance, arguably today's most-notable (and indisputably the most-noted) political-economic phenomenon, refers directly to the interdependence of domestic politics, policies, and policymakers. International economic integration is widely considered a root cause of the recent cross-national spread of economic liberalization and the so-called *Washington Consensus*, and many commentators even see international waves of partisan governments and votes as resulting from some interdependence in mass opinion and vote choices (but *cf.* Kayser 2007).

The substantive/theoretical ubiquity and centrality of interdependence across political-science discrete-choice contexts notwithstanding, studies that accord interdependence explicit attention are uncommon. The rare exceptions include the diffusion and network literatures cited above; Ward, Gleditsch, and colleagues⁴ and Signorino and coauthors⁵ in international relations; Li & Thompson (1975), Govea & West (1981), and Brinks & Coppedge (2006) on coups, riots, and revolutions, respectively; Schofield et al. (2003) on citizens' votes and Lacombe & Shaughnessy (2005) on legislators' votes; and Mukherjee & Singer (2007) on inflation targeting.

⁴ Shin & Ward 1999, Gleditsch & Ward 2000, Gleditsch 2002, Ward & Gleditsch 2002, Hoff & Ward 2004, Gartzke & Gleditsch 2006, Salehyan & Gleditsch 2006, Gleditsch 2007.

⁵ Signorino 1999, 2002, 2003, Signorino & Yilmaz 2003, Signorino & Tarar 2006.

Likewise, despite the manifest interdependence in social-science discrete-choices, assumptions of *in*dependence pervade almost all empirical analyses of them, even in those research areas that give interdependence greater substantive and theoretical weight. Empirical models of war in which the dependence of one state's choices on those of others enters explicitly are rare.⁶ Empirical models of policy, institution, or regime diffusion often do account interdependence explicitly by including as explanators (weighted) averages or sums of other units' outcomes (e.g., the number of other states that have adopted a policy or treaty), but the endogeneity of this *spatial lag* is rarely confronted. Typically, diffusion researchers time-lag these spatial lags, as in the sophisticated event-history analyses of modern applications for example, and this *can* suffice to evade the simultaneity bias (see, e.g., Beck et al. 2006), but only *if* and so far as (i) actual interdependence transpires only with a lag, (ii) with actual lag periodicity and lag structure identical to that of the empirical observations and specification, and (iii) that the empirical model of spatiotemporal dynamics is adequate to prevent the past bleeding into present through mismeasurement/misspecification.⁷ Even if this time-lag strategy sufficed on these grounds, moreover, another methodological problem arises for current practice in binary-outcome models. With binary outcomes, placing the actual policies of other units (or their weighted sums or averages) on the right-hand side (as opposed to the latent variables) typically produces a kind of logical inconsistency (as explained later; see Heckman 1978).

Likewise, empirical network analyses, including the most recent and exciting contributions in random-graph (e.g., Robins et al. 2007; Hunter et al. 2008), longitudinal-network (Snijders 2005), and/or network-coevolution (e.g., Snijders et al. 2007) modeling, fail to address fully and/or directly the interdependence of their binary outcomes. Random-graph models proceed, instead, by modeling summary statistics for the entire network of binary ties rather than modeling specific ties, thereby

⁶ Ward, Gleditsch, and colleagues and the Signorino and coauthors are among the few exceptions (see note 4).

⁷ As Beck et al. (2006), e.g., note, adequacy of the spatiotemporal dynamic model can and should be tested. We have not seen these tests conducted in the diffusion literature though, nor, usually, sign that researchers are aware of the issue.

evading any direct model of the (inter)dependence of specific ties i to j (with) on specific k to l ties. Rather, one models the implications for some network summary-statistic(s) of particular behavioral tendencies, some of which might implicate dependencies: e.g., a tendency toward transitive-triplets (A-C and A-B ties increase the likelihood of B-C ties). Longitudinal-network models generally apply temporal-sequencing strategies similar in essence to those of the diffusion literatures. Similarly, network-behavior coevolution models assume the network-tie decisions and behavioral choices of units are independent of each other and of other units' decisions and choices, conditional upon the existing network and set of units' characteristics (i.e., the pre-existing set of network and behavioral choices). In sum, none of these approaches allows a direct, simultaneous interdependence of binary outcomes (and, recall, *simultaneous* means within observational period, as effectively modeled).⁸

Working under the incorrect assumption of independence, of course, threatens over-confidence or inefficiency in the best of circumstances, and usually bias and inconsistency as well. Inclusion of spatial lags to reflect interdependence would seem advisable, but such lags are endogenous and so introduce simultaneity biases.⁹ For the linear-regression case, we have argued and shown elsewhere (2003, 2004ab, 2005ab, 2006bc, 2007bcd, 2008ab, 2009ab) that serious omitted-variable biases arise when spatial lags are excluded in the presence of interdependence and that redressing this issue by explicit inclusion of spatial lags to reflect interdependence is generally of first-order benefit relative to the problems induced by spatial-lag endogeneity. However, these simultaneity biases do become appreciable as interdependence strengthens, so we also covered in these previous works methods for gauging that strength, for redressing the simultaneity issues of spatial lags, and for calculating and

⁸ We would also note in passing here, although this is an argument for fuller development in another venue, that the presence of interdependence generally violates the crucial SUTVA assumption of matching methods for causal inference, invalidating, or at least seriously complicating, such less structural/parametric approaches to empirical inference.

⁹ The inclusion of weighted sums of other units' outcomes also introduces measurement error insofar as interdependence truly arises through *expectations* of other units' outcomes. Substantively, alternative interdependence mechanisms may suggest diffusion either of outcomes or expected-outcomes, but only the latter mechanism can be identified logically in the case of binary outcomes. Again, see further explanation below (and original exposition in Heckman 1978).

presenting estimates of spatially/spatiotemporally dynamic effects and their certainty, but (almost) exclusively in the linear-regression context (2007d briefly discussed spatial probit). This paper starts a similar exploration of spatial models of binary outcomes, and more generally of qualitative and limited dependent variables (henceforth: *QualDep*), where the substantive and theoretical importance of interdependence, the empirical problems created by its omission, and the methodological challenges raised by the endogeneity of its appropriately explicit inclusion are all even greater.

II. The Econometric Problem

Methods for properly estimating and analyzing models of interdependent qualitative or limited dependent variables (henceforth: *QualDep* models) have received significant attention in the spatial-econometric literature recently. Most of this research considers the spatial-probit model¹⁰ with interdependence in the latent-variable, i.e., in the unobserved argument to the probit-modeled probability of a binary outcome. Models of spatial sample-selection (spatial Tobit or Heckit: McMillen 1995, Smith & LeSage 2004, Flores-Lagunes & Schnier 2006), spatial multinomial-probit (McMillen 1995, Bolduc et al. 1997), and spatial discrete-duration (Phaneuf & Palmquist 2003), all of which closely resemble the spatial probit, have also been suggested, as have models of interdependent survival (Hays & Kachi 2008) or of survival with spatial frailty (Banerjee et al. 2004, Darmofal 2007) and of spatial counts (Bhati 2005), including a zero-inflated-count model (Rathbun & Fei 2006). Spatial probit is far the most-common *S-QualDep* model in applied research, however.¹¹

Several estimation strategies have been suggested for the spatial-probit model. McMillen (1992) first suggested an EM algorithm, which rendered the spatial-probit's non-additively-separable log-likelihood (see below) estimable for the first time, but the strategy also did not provide standard-

¹⁰ Spatial logit has also been suggested (e.g., Dubin 1997; Lin 2003; Autant-Bernard 2006), but spatial probit dominates the methodological and applied literatures, likely due to the relatively greater feasibility of working with n -dimensional normal (as opposed to extreme-value) distributions as necessary to incorporate the interdependence directly.

¹¹ E.g., Holloway et al 2002, Beron et al 2003, Coughlin et al 2003, Murdoch et al 2003, Novo2003, Schofield et al 2003, Garrett et al. 2005, Lacombe & Shaughnessy 2005, Autant-Bernard 2006, Rathbun&Fei 2006, Mukherjee&Singer 2007.

errors for the crucial spatial-dependence parameter and required arbitrary parameterization of the heteroscedasticity that dependence induces (see below). McMillen (1995) and Bolduc et al. (1997) applied simulated-likelihood strategies to estimate their spatial-multinomial-probit models, and Beron et al. (2003) and Beron & Vijverberg (2004) advanced a recursive-importance-sampling (RIS) estimator in that line. LeSage (1999, 2000, 2004) introduced a Bayesian strategy of Markov-Chain-Monte-Carlo (MCMC) by Metropolis-Hastings-within-Gibbs sampling. Fleming (2004) reviews these two families and simpler, if approximate, strategies allowing spatial interdependence in linear or nonlinear probability models¹² estimable by nonlinear least-squares, generalized linear-models, or generalized linear-mixed-models. Pinkse & Slade's (1998) two-step GMM estimator has also seen some use in the literature, but the RIS and Bayesian strategies have dominated recent applications.

The remainder of this section considers the spatial-probit model and RIS and Bayesian strategies for estimating it. The structural model for the latent variable of the spatial probit takes the form:

$$\mathbf{y}^* = \rho \mathbf{W} \mathbf{y}^* + \mathbf{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (1),$$

which can be written in reduced form as:

$$\mathbf{y}^* = (\mathbf{I} - \rho \mathbf{W})^{-1} \mathbf{X} \boldsymbol{\beta} + \mathbf{u}, \quad \text{with } \mathbf{u} = (\mathbf{I} - \rho \mathbf{W})^{-1} \boldsymbol{\varepsilon} \quad (2),$$

Latent-variable \mathbf{y}^* links to the observed binary-outcome, \mathbf{y} , through the measurement equation:

$$y_i = \{1 \text{ if } y_i^* > 0 ; 0 \text{ if } y_i^* \leq 0\} \quad (3).$$

The probabilities that the i^{th} observation is one are calculated as follows:

$$\begin{aligned} p(y_i = 1 | \mathbf{X}) &= p\left(\left[(\mathbf{I} - \rho \mathbf{W})^{-1} \mathbf{X} \boldsymbol{\beta}\right]_i + \left[(\mathbf{I} - \rho \mathbf{W})^{-1} \boldsymbol{\varepsilon}\right]_i > 0\right) \\ &= p\left(u_i < \left[(\mathbf{I} - \rho \mathbf{W})^{-1} \mathbf{X} \boldsymbol{\beta}\right]_i / \sigma_i\right) = \Phi_i \left\{ \left[(\mathbf{I} - \rho \mathbf{W})^{-1} \mathbf{X} \boldsymbol{\beta}\right]_i / \sigma_i \right\} \end{aligned} \quad (4).^{13}$$

Thus, as in the standard probit, a cumulative-normal distribution, $\Phi\{\cdot\}$, gives the probability that

¹² Even the linear-probability model becomes nonlinear in parameters given the spatial multiplier, $(\mathbf{I} - \rho \mathbf{W})^{-1}$.

¹³ In the middle step, note that the symmetry about zero of $\boldsymbol{\varepsilon}$, and so of \mathbf{u} , implies that $p(-u_i < x) = p(u_i < x)$ for any x .

the systematic component, $[(\mathbf{I} - \rho\mathbf{W})^{-1}\mathbf{X}\boldsymbol{\beta}]_i / \sigma_i$, exceeds the stochastic component, u_i . However, in spatial probit, the interdependence of the y_i^* induces a non-sphericity of the stochastic components; specifically, \mathbf{u} is distributed n -dimensional *multivariate* normal with variance-covariance matrix $[(\mathbf{I} - \rho\mathbf{W})'(\mathbf{I} - \rho\mathbf{W})]^{-1}$ (and mean $\mathbf{0}$). Intuitively, $\boldsymbol{\varepsilon}$ is *multivariate* normal with mean $\mathbf{0}$ and spherical variance-covariance $\sigma^2\mathbf{I}$, with σ^2 normalized to 1 as usual for a probit model; therefore:

$$\begin{aligned} V[(\mathbf{I} - \rho\mathbf{W})^{-1}\boldsymbol{\varepsilon}] &= [(\mathbf{I} - \rho\mathbf{W})^{-1}]'V(\boldsymbol{\varepsilon})[(\mathbf{I} - \rho\mathbf{W})^{-1}] \\ &= [(\mathbf{I} - \rho\mathbf{W})^{-1}]'\mathbf{I}[(\mathbf{I} - \rho\mathbf{W})^{-1}] = [(\mathbf{I} - \rho\mathbf{W})'(\mathbf{I} - \rho\mathbf{W})]^{-1} \end{aligned} \quad (5).$$

The probability that $[(\mathbf{I} - \rho\mathbf{W})^{-1}\mathbf{X}\boldsymbol{\beta}]_i / \sigma_i$ exceeds u_i is read from the i^{th} marginal distribution of this multivariate cumulative-normal, denoted $\Phi_i\{\cdot\}$, which requires integrating that joint distribution over the other $n-1$ dimensions. In (4), σ_i^2 is the ii^{th} element of variance-covariance (5), which is not a constant (set to one) as in standard-probit. I.e., spatial interdependence induces heteroscedasticity. This heteroscedasticity and, more fundamentally, the interdependence (i.e., the *non*-independence) of the u_i , render standard probit inappropriate. Because the outcomes are interdependent, their joint distribution is not the product of the n marginal distributions, so one does not maximize sums of logs of n additively separable one-dimensional probabilities. They are interdependent, so one maximizes the log of one non-separable n -dimensional distribution. Finally, notice also that the i^{th} observation probability depends on the entire matrix \mathbf{X} and vector $\boldsymbol{\varepsilon}$. This follows from the nonlinearity of the sigmoidal probit function, which implies that effects depend on where along the S -shape they occur, and where on that S -curve one lies depends on all of \mathbf{X} and $\boldsymbol{\varepsilon}$ given the dependence of y_i^* on $\mathbf{W}\mathbf{y}^*$.

The spatial-error version of the probit model is slightly simpler, taking the form:

$$\mathbf{y}^* = \mathbf{X}\boldsymbol{\beta} + \mathbf{u} \quad (6),$$

with $\mathbf{u} = (\mathbf{I} - \lambda\mathbf{W})^{-1}\boldsymbol{\varepsilon}$, and having the marginal probabilities:

$$p(y_i = 1 | \mathbf{x}_i) = p(u_i < \mathbf{x}_i \boldsymbol{\beta} / \sigma_i) = \Phi_i \{ \mathbf{x}_i \boldsymbol{\beta} / \sigma_i \} \quad (7),$$

where \mathbf{x}_i is the i^{th} row of \mathbf{X} . Again, these u_i are heteroskedastic and the probability derives from the i^{th} marginal distribution of a multivariate cumulative-normal with means $\mathbf{0}$ and variance-covariance $[(\mathbf{I} - \lambda \mathbf{W})'(\mathbf{I} - \lambda \mathbf{W})]^{-1}$, so spatial-error probit models entail the same estimation and interpretation complications as spatial-lag models. (Mixed spatial-lag/spatial-error models are also possible, but they have received little attention.) In the spatial-error model, because the interdependence operates only through $\boldsymbol{\varepsilon}$ and not all of \mathbf{y}^* , the position of the i^{th} observation on the sigmoidal probit-function depends on the entire vector $\boldsymbol{\varepsilon}$ but only on that observation's independent-variable values, \mathbf{x}_i .

Special circumstances might allow standard-probit estimation of spatial-lag models, but we view these as highly atypical. For instance, Anselin (2006) notes that, in the conditional counterpart of (1),

$$y_i^* = \rho \sum_j w_{ij} E(y_j^* | \mathbf{X}) + \mathbf{x}_i \boldsymbol{\beta} + \boldsymbol{\varepsilon} \quad (8),$$

$E(y_j^* | \mathbf{X})$ could be estimated by $\sum_j w_{ij} y_j$, the spatially weighted average of actual outcomes in units j .¹⁴ However, this spatial lag could be included as a regressor without introducing endogeneity problems only under stringent conditions that ensure other units' observations j are not jointly determined with those of i , and that "coding methods ensure that the sample does not contain these neighbors" (Anselin 2006). This means that any units j from which diffusion to any i in the sample is non-negligible (at any order spatial-lag) must be excluded from the sample but used in constructing the $\mathbf{W} \mathbf{y}$ spatial lag for the retained observations i . Alternatively, all i 's neighboring j according to \mathbf{W} must be exogenous to i for all i in the sample; i.e., feedback must be directional and orderable from j 's to i 's only, severing feedback from i back to itself. Moreover, while some substantive-theoretical contexts might suggest that interdependence propagates through the actual outcome rather

¹⁴ Note that the row-normalization here means that the resulting seeming replication of the summed weights is irrelevant.

than the latent variable, such a model is not generally possible because, indirectly via feedback, y_i would generate y_i^* but also, directly, y_i is generated by $\Phi(y_i^*)$.¹⁵ Conditions like those described above allow direct inclusion of $\mathbf{W}y$ because they sever such indirect generation of y_i^* by y_i . These limitations are usually prohibitive practically, though contexts where such directional ordering and such omissions of certain j may be defensible are imaginable. Swank (2006, 2007), e.g., argues that U.S. tax policies exclusively lead others' tax policies, and he excludes all U.S. data in his tax-competition empirics, reserving those U.S. data solely for the role of spatial lag. If valid, these arguments and sample-exclusions would allow standard-probit estimation.

We focus on the unconditional, simultaneous spatial-lag (and, less so, spatial-error) model in the rest of the paper. We ignore the conditional model as it will usually be inapplicable and, anyway, raises fewer interpretation and no estimation complications. Similarly, we will not discuss the time-lagged spatial-lag model further because the conditions discussed above for the practical adequacy of the strategy seem restrictive for social-science applications and because, even if otherwise adequate, the strategy will not generally evade the complications of lagged binary-dependent-variables.¹⁶ We also do not discuss tests of the adequacy of time-lagged spatial-lag models or specification tests of spatial-lag vs. spatial-error vs. non-spatial models here, though these tests are important to consider, especially given the complexity and computational intensity of valid estimation strategies for full,

¹⁵ The requirement applies to any simultaneous feedback among endogenous qualitative variables, as perhaps first noted by Heckman (1978) in the context of a system of 2 endogenous equations, at least one of them being qualitative and modeled by a latent variable crossing a threshold. He states: "A necessary and sufficient condition for [sensitivity of such a system of endogenous latent-variable equations is] that the probability of the event $d_i=1$ is not a determinant of the event... [This] principal assumption essentially requires that the latent variable y^* and not the measured variable y appears [on the right-hand side of the] structural equation" (pp. 936-7). The same limitation does not quite obtain for temporal dependence, however. Since time is unidirectional, one may be able to rely on pre-determinedness of y_{t-1} , i.e., the indirect feedback from y_t to y_{t-1} does not occur (given sufficiently full and accurate specification of the temporal dynamics). Still, conditions for proper identification of just a temporally dynamic model with lagged binary-dependent-variables remain far from straightforward (see, e.g., Chamberlain 1993, Honore & Kyriazidou 2000).

¹⁶ Monte Carlo simulation exploring the sensitivity of the time-lagging spatial-dependence strategy to validity of the lagged-interdependence-only assumption, to the periodicity-matching assumption, and to the empirical adequacy of the spatiotemporal dynamic model and tests thereof are important analyses that remain for the future.

simultaneous spatial probit. We refer the reader to Pinkse & Slade (1998), Pinkse (1999), Kelejian & Prucha (2001), and, for a recent review, Anselin (2006). Our considerations focus on spatial-probit estimation by RIS and by Bayesian MCMC methods, and their comparison to standard probit estimation with the endogenous spatial-lag, W_y , erroneously included as a regressor, which is current standard-practice in empirical work where interdependence of binary outcomes is addressed.

III. The RIS and Bayesian Estimators for Spatial Probit

LeSage (1999, 2000) suggests using Bayesian Markov-Chain-Monte-Carlo (MCMC) methods to surmount the estimation complications introduced by the n -dimensional cumulative-normal in the spatial-probit likelihood (posterior). The basic idea of Monte Carlo (simulation) methods is simple:¹⁷ if one can characterize the joint distribution (likelihood or posterior) of the quantities of interest (parameters), then one can simply sample (take random draws) from that distribution and calculate the desired statistics in those samples. With sufficient draws, the sample statistics can approximate the population parameters¹⁸ they aim to estimate arbitrarily closely. In basic Monte-Carlo simulation, the draws are independent and the target distribution is specified directly. In MCMC, each draw is dependent on the previous one in a manner that generates samples with properties mirroring those of the joint population, using just the conditional distribution of each parameter. This is useful where the joint distribution is not expressible directly or, as with spatial probit, where its complexity makes direct sampling from the joint distribution prohibitively difficult and/or time-consuming.

We can describe Gibbs sampling, the simplest and most-common of the MCMC family, thusly: Given distributions for each parameter conditional on the other parameters, one can cycle through draws from those conditional distributions, eventually reaching a *convergent* state past which point

¹⁷ Our simple introduction draws heavily from Gill's (2002) wonderful text on Bayesian methods.

¹⁸ Recall that the "population parameters" that can be arbitrarily closely approximated will usually be some *estimates* in an application, like the spatial-probit parameter-estimates, not the "true parameters" of course (which latter concept is somewhat awkward in Bayesian terminology anyway).

all subsequent draws will be from the targeted posterior joint-distribution. To elaborate: first express the distribution for each parameter conditional on all the others, then choose (arbitrary) starting values for those parameters and draw a new value for the first parameter conditional on the others' starting values. Then, conditional on this new draw of the first parameter and starting values for the rest, draw a new value for the second parameter from its conditional distribution. Continue thusly until all parameters have their first set of drawn values, then return to the first parameter and draw its second simulated value conditional on the others' first draws. Cycle thusly for some large number of iterations, and, under rather general conditions, the limiting (asymptotic) distribution of this set of parameter draws is the desired joint posterior-distribution. Thus, after having gathered some very large set of parameter-vector values by this process, discard some large initial set of draws (the *burn-in*) and base inferences on sample statistics from the remaining set of parameter vectors. A typical burn-in might be 1000 draws, and inferences might be based on the next 5000 or 10,000. Also, since each draw is conditional on the previous one's drawn values, autocorrelation typically remains, so "thinning" the post-burn-in sample by using every, say, third or fifth draw may boost efficiency.

The drawbacks of MCMC may be obvious from what we have said and declined to say. First, no universal tests exist to verify that *convergence* has occurred, so a burn-in may appear sufficient in that the next 5000 drawn parameter-vectors seem to follow some circumscribed bounds and behavior of some unknown target distribution (i.e., the sampler may seem to have *settled down*) only to have the 5001st leap into a new range and proceed toward convergence elsewhere. Second, despite their Markov-Chain origins, adjacent draws are asymptotically serially uncorrelated, but this *asymptopia* may not arrive within practical limits, and thinning may be insufficient help or too computationally costly. Third, the starting values are likewise asymptotically irrelevant, assuming the supplied set of conditional distributions properly could come from a valid joint distribution, but, as the previous two

caveats imply, starting values may matter short of convergence, arrival at which is not verifiable.¹⁹ These issues concern careful researchers, and many diagnostics and tests for non-convergence, serial correlation, or starting-value sensitivity, and some strategies for ameliorating them, exist (all imperfect, but useful still). However, the concerns do not outweigh the remarkably flexible utility of the Gibbs sampler, either in general or specifically in its application to spatial-probit estimation.

All but one of the conditional distributions for the spatial-probit-model parameters (given below) are standard, so the Gibbs sampler is useful for them. The crucial spatial-lag-coefficient, ρ , has the lone non-standard conditional-distribution; for it, Metropolis-Hastings sampling is used. Metropolis-Hastings differs from Gibbs sampling in the former's *seeding* or *jump* distribution from which values are drawn and then accepted or rejected as the next sampled parameters, depending on how they compare to a suitably transformed expression of the target distribution.²⁰ The Bayesian spatial-probit estimator (LeSage 1999, 2000) uses Metropolis-Hastings for ρ within the Gibbs sampler procedure for the other parameters. Of course, this step adds some to the estimator's computational intensity.

With this brief introduction to Bayesian MCMC estimation by Gibbs and Metropolis-Hastings sampling, we now introduce their application to the spatial-probit model. We follow LeSage (2000) to state the likelihood in terms of the latent outcome, \mathbf{y}^* —an additional conditional distribution will later apply (3) to convert unobserved \mathbf{y}^* to observed \mathbf{y}^{21} —for the spatial-lag model (1) as:

$$L(\mathbf{y}^*, \mathbf{W} | \rho, \boldsymbol{\beta}, \sigma^2) = \frac{1}{2\pi\sigma^{2(n/2)}} |\mathbf{I}_n - \rho\mathbf{W}| e^{-\frac{1}{2\sigma^2}(\boldsymbol{\epsilon}'\boldsymbol{\epsilon})} \quad (9),$$

¹⁹ The conditional distributions must also be expressible and sufficiently tractable to make so many draws a practicality.

²⁰ To elaborate: to sample from some non-standard density $f(\cdot)$, let x_0 be the current draw from $f(\cdot)$, beginning with an arbitrary starting value. Consider a candidate next value, x_1 , for x given by $x_1 = x_0 + cZ$ with Z being drawn from a standard-normal distribution and c a given constant. Then, we assign a probability of accepting this candidate as the next value of x in our MCMC as $p = \min\{1, f(x_1)/f(x_0)\}$. I.e., we draw from a Uniform(0,1) distribution, and, if $U < p$, the candidate x_1 becomes the next x ; if $U > p$ the next x remains x_0 . Metropolis-Hastings is thus one type of *rejection sampling*.

²¹ This stratagem also enables LeSage to express the spatial-Tobit model by this same likelihood, adding a conditional distribution later to generate latent variables z for censored observations instead of one to generate $y = (0, 1)$ for the probit.

where $\boldsymbol{\varepsilon} = (\mathbf{I}_n - \rho \mathbf{W}) \mathbf{y}^* - \mathbf{X} \boldsymbol{\beta}$. The likelihood for spatial-error probit model (6) is the same but with $\boldsymbol{\varepsilon} = (\mathbf{I}_n - \rho \mathbf{W}) (\mathbf{y}^* - \mathbf{X} \boldsymbol{\beta})$, where ρ serves here for λ in (6). Diffuse priors yield joint posterior-density:

$$p(\rho, \boldsymbol{\beta}, \sigma | \mathbf{y}^*, \mathbf{W}) \propto |\mathbf{I}_n - \rho \mathbf{W}| \sigma^{-(n+1)} e^{-\frac{1}{2\sigma^2}(\boldsymbol{\varepsilon}'\boldsymbol{\varepsilon})} \quad (10).$$

One can now derive the conditional posterior densities for ρ , $\boldsymbol{\beta}$, and σ for the sampler. First:

$$p(\sigma | \rho, \boldsymbol{\beta}) \propto \sigma^{-(n+1)} e^{-\frac{1}{2\sigma^2}(\boldsymbol{\varepsilon}'\boldsymbol{\varepsilon})} \quad (11).$$

Notice that conditioning on ρ allows $|\mathbf{I}_n - \rho \mathbf{W}|$ to be subsumed into the constant of proportionality and that (11) implies $\sigma^2 \sim \chi_n^2$, a standard distribution facilitating the Gibbs sampler. Next,

$$p(\boldsymbol{\beta} | \rho, \sigma) \sim N[\tilde{\boldsymbol{\beta}}, \sigma_\varepsilon^2 (\mathbf{X}' \mathbf{C}' \mathbf{C} \mathbf{X})^{-1}] \quad (12),$$

where, in spatial lag, $\mathbf{C} = \mathbf{I}_n$ and $\tilde{\boldsymbol{\beta}} = (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' (\mathbf{I}_n - \rho \mathbf{W}) \mathbf{y}^*$, and, in spatial error, $\mathbf{C} = (\mathbf{I}_n - \rho \mathbf{W})$ and $\tilde{\boldsymbol{\beta}} = (\mathbf{X}' \mathbf{C}' \mathbf{C} \mathbf{X})^{-1} \mathbf{X}' \mathbf{C}' \mathbf{C} \mathbf{y}^*$. The conditional multivariate-normality of $\boldsymbol{\beta}$ allows the Gibbs sampler for it also, but ρ has non-standard conditional distribution, requiring Metropolis-Hastings sampling:

$$p(\rho | \boldsymbol{\beta}, \sigma) \propto |\mathbf{I}_n - \rho \mathbf{W}| \sigma^{-(n+1)} e^{-\frac{1}{2\sigma^2}(\boldsymbol{\varepsilon}'\boldsymbol{\varepsilon})} \quad (13),$$

with $\boldsymbol{\varepsilon}$ defined as given above for the spatial-error and the spatial-lag models.²² Finally, we add the conditional distribution, namely a truncated normal, that translates \mathbf{y}^* to \mathbf{y} :

$$f(z_i | \rho, \boldsymbol{\beta}, \sigma) \sim N(\tilde{y}_i, \sigma_i^2), \text{ left- or right-truncated at 0 as } y_i = 1 \text{ or } 0 \quad (14),$$

where \tilde{y}_i is the predicted value of y_i^* (the i^{th} element of $(\mathbf{I}_n - \rho \mathbf{W})^{-1} \mathbf{X} \boldsymbol{\beta}$ for spatial-lag or of $\mathbf{X} \boldsymbol{\beta}$ for spatial-error models) and the variance of \tilde{y}_i is $\sum_j \omega_{ij}^2$ with ω_{ij} the i^{th} element of $(\mathbf{I}_n - \rho \mathbf{W})^{-1} \boldsymbol{\varepsilon}$.²³

²² Anselin (1988) shows that the minimum and maximum eigenvalues of a standardized spatial-weight matrix, \mathbf{W} , bound ρ to $1/\lambda_{\min} < \rho < 1/\lambda_{\max}$. One could add this constraint to the rejection sampling, but our preliminary simulations seem to suggest that the model-estimates have better properties if one instead applies the wider bounds of $(-1, 1)$.

²³ Spatial Tobit replaces (14) with: $f(z_i | \rho, \boldsymbol{\beta}, \sigma) \sim \begin{cases} [1 - \Phi(y_i^* / \sigma_i)]^{-1} \exp[-(z_i - y_i^*)^2 / 2\sigma_i], & \text{if } z_i \leq 0. \\ 0, & \text{if } z_i > 0 \end{cases}$. The Tobit allows

With all the conditional distributions, we can implement MCMC to estimate the model thus:²⁴

1. Use expression (11) to draw σ_1 using starting values $\rho_0, \boldsymbol{\beta}_0, \sigma_0$.
2. Use σ_1, ρ_0 , and expression (12) to draw $\boldsymbol{\beta}_1$.
3. Use $\sigma_1, \boldsymbol{\beta}_1$, and expression (13) to draw ρ_1 by Metropolis-Hastings sampling.
4. Sample the outcome, z_i , using the censoring distribution (14) and $\sigma_1, \boldsymbol{\beta}_1$, and ρ_1 .
5. Return to step 1 incrementing the subscript counters by one.

After a sufficient burn-in—our simulation and application experiences so far suggest at least 1000 is advisable—the distributions of $\sigma_1, \boldsymbol{\beta}_1$, and ρ_1 will have reached convergence and subsequent draws on the parameters may be used to give their estimates (as means or medians of some large number of draws) and estimates of their certainty (as standard deviations or percentile ranges).²⁵

A frequentist approach has also been suggested, Recursive Importance-Sampling (RIS), which also uses simulation to approximate probabilities difficult to calculate analytically, for estimating spatial-lag or spatial-error probit. We introduce RIS following Vijverberg's (1997) notation. To approximate an n -dimensional cumulative multivariate-normal distribution, e.g.,

$$p = \int_{-\infty}^{\mathbf{x}_0} f_n(\mathbf{x}) d\mathbf{x}, \quad (15),$$

where $f_n(\mathbf{x})$ is the density and $[-\infty, \mathbf{x}_0]$ the interval over which we want to integrate, we first choose a n -dimensional sampling-distribution with well-known properties and label a truncation of

$\sigma_i^2 = \sigma_{\varepsilon_i}^2 \sum_j \omega_{ij}^2$, but the probit must scale $\sigma_{\varepsilon_i}^2$ to 1, it and $\boldsymbol{\beta}$ not being separately identified for binary-outcome models.

²⁴ In assigning diffuse priors to the parameters, LeSage (2000) also relaxes the assumption of homoskedasticity in $\boldsymbol{\varepsilon}$, allowing $V(\boldsymbol{\varepsilon})$ to vary arbitrarily by observation i . This allows exploration of variation in model fit and identification of and robustness to potential outliers, but creates as many parameters to estimate as observations. LeSage circumvents that issue by specifying an informative prior for those relative-variance parameters, specifically one suggested by Geweke (1993) that, *inter alia*, has the useful property of yielding a distribution of $\boldsymbol{\varepsilon}$ consistent with a probit choice-model as the Gewekian-distribution parameter, q , goes to infinity, and that at $q \approx 7.5$ yields a choice-model approximating logit. The posterior-estimates of q , may therefore be used to test logit versus probit (versus un-named possibilities $q \neq 7.5$ and $q \neq \infty$).

Allowing arbitrary relative-variance requires the additional (informative) Gewekian prior and a (diffuse) hyper-prior on its parameter, q ; produces more complicated expressions for the conditional distributions of $\sigma, \rho, \boldsymbol{\beta}$; and adds a conditional distribution (fortunately standard: χ^2_{q+1}) for the relative variances, v_i . The steps below would now also include conditioning on starting values for, and then the previous draws of, \mathbf{v} , and a step inserted between 2 and 3 would draw the next \mathbf{v} from χ^2_{q+1} conditional on the current $\sigma, \rho, \boldsymbol{\beta}$. Notice that setting the hyper-prior for q determinately to a large number (or 7.5) yields spatial probit (or logit) without heteroscedasticity/outlier-robustness.

²⁵ Thinning may also be advisable, although we have not yet explored the possibility.

this sampling distribution with support over $[-\infty, \mathbf{x}_0]$ the *importance distribution*. Defining $g_n^c(\mathbf{x})$ as the density for this n -dimensional importance distribution, we then multiply and divide the right-hand-side of the integral we wish to calculate, (15), by this density, which simply restates (15) as:

$$p = \int_{-\infty}^{\mathbf{x}_0} \frac{f_n(\mathbf{x})}{g_n^c(\mathbf{x})} g_n^c(\mathbf{x}) d\mathbf{x} \quad (16).$$

By definition, the solution to this integral is a mean because $g_n^c(\mathbf{x})$ is a valid *pdf* over the integral's range, so (16) gives the probability sought, p , as the mean of $f_n(\mathbf{x})/g_n^c(\mathbf{x})$, which we can estimate using a sample of R draws of the $n \times 1$ vector \mathbf{x} from the importance distribution. Formally:

$$p = E \left[\frac{f_n(\mathbf{x})}{g_n^c(\mathbf{x})} \right] \approx \frac{1}{R} \sum_{r=1}^R \frac{f_n(\tilde{\mathbf{x}}_r)}{g_n^c(\tilde{\mathbf{x}}_r)} \equiv \hat{p} \quad (17).$$

To implement the RIS estimator, we simply draw \mathbf{x} from the importance-distribution, for which we will use a truncated multivariate (independent) normal,²⁶ and calculate $f_n(\mathbf{x})/g_n^c(\mathbf{x})$.

Again, in the standard probit-model with independent errors, the numerator would simply sum n univariate cumulative standard-normal distributions, which is manageable. In spatial probit, with its interdependent errors, however, the numerator is a single n -dimensional cumulative-normal:

$$p(\mathbf{u} < \mathbf{v}) \quad (18),$$

with \mathbf{u} the $n \times 1$ vector of errors distributed $MVN(\mathbf{0}, \Sigma)$ and $\Sigma = (\mathbf{I} - \rho \mathbf{W})' (\mathbf{I} - \rho \mathbf{W})^{-1}$, and with \mathbf{v} an $n \times 1$ vector $\mathbf{v} = \mathbf{Q}(\mathbf{I} - \rho \mathbf{W})^{-1} \mathbf{X}\beta$, where \mathbf{Q} is a diagonal matrix with diagonals $q_i = 2y_i - 1$.²⁷ The RIS estimator for spatial probit exploits the fact that, as a variance-covariance matrix, Σ is positive definite, so a Cholesky decomposition exists such that $\Sigma^{-1} = \mathbf{A}'\mathbf{A}$, with \mathbf{A} being an upper-triangular

²⁶ Other importance distributions, such as a t or a uniform may be used. With a normal importance-distribution, RIS is equivalent to the better-known GHK (Geweke-Hajivassiliou-Keane) simulation estimator.

²⁷ Note that $q=2y_i-1$ is 1 for $y_i=1$ and -1 for $y_i=0$; thus, multiplying by \mathbf{Q} serves to select the right sign on the systematic component up to which to integrate the distribution of the stochastic component \mathbf{u} . See, e.g., Greene (2008:778).

matrix and $\boldsymbol{\eta} = \mathbf{A}\mathbf{u}$ giving n independent standard-normal variables, $\boldsymbol{\eta}$. (This exploitation is familiar as the same one applied in GLS.) Let $\mathbf{B} \equiv \mathbf{A}^{-1}$; substituting $\mathbf{u} = \mathbf{A}^{-1}\boldsymbol{\eta} \equiv \mathbf{B}\boldsymbol{\eta}$ into (18) then gives:

$$\Pr(\mathbf{B}\boldsymbol{\eta} < \mathbf{v}) = \Pr \left(\begin{bmatrix} b_{1,1} & b_{1,2} & \cdots & \cdots & b_{1,n} \\ 0 & b_{2,2} & \ddots & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & b_{n-1,n-1} & b_{n-1,n} \\ 0 & \cdots & 0 & 0 & b_{n,n} \end{bmatrix} \begin{bmatrix} \eta_1 \\ \vdots \\ \vdots \\ \eta_{n-1} \\ \eta_n \end{bmatrix} < \begin{bmatrix} v_1 \\ \vdots \\ \vdots \\ v_{n-1} \\ v_n \end{bmatrix} \right) \quad (19).$$

The elements of the $n \times 1$ vector $\boldsymbol{\eta}$ are independent, so the probability in (19) can be calculated by first evaluating the cumulative-normal distribution function at the implied upper bounds, which are determined recursively starting with the last observation, and then multiplying these probabilities. To determine these upper bounds, start by solving the inequalities in (19) for the vector $\boldsymbol{\eta}$:

$$\Pr \left(\begin{bmatrix} \sum_{i=1}^n b_{1,i} \eta_i \\ \vdots \\ b_{n-1,n-1} \eta_{n-1} + b_{n-1,n} \eta_n \\ b_{n,n} \eta_n \end{bmatrix} < \begin{bmatrix} v_1 \\ \vdots \\ v_{n-1} \\ v_n \end{bmatrix} \right) = \Pr \left(\begin{bmatrix} \eta_1 \\ \vdots \\ \eta_{n-1} \\ \eta_n \end{bmatrix} < \begin{bmatrix} b_{1,1}^{-1} \left(v_1 - \sum_{i=2}^n b_{1,i} \eta_i \right) \\ \vdots \\ \vdots \\ b_{n-1,n-1}^{-1} \left(v_{n-1} - b_{n-1,n} \eta_n \right) \\ b_{n,n}^{-1} v_n \end{bmatrix} \right) \quad (20)$$

Next, calculate the upper bound for the truncated-normal distribution of the n^{th} observation, which is $b_{n,n}^{-1} v_n$. Call the cumulative standard-normal evaluated at this upper bound p_n . Then take a draw from the standard-normal distribution truncated at $b_{n,n}^{-1} v_n$; call that draw $\tilde{\eta}_n$ and use it to calculate the upper bound for the truncated-normal distribution for the $(n-1)^{\text{th}}$ observation conditional on the n^{th} as $b_{n-1,n-1}^{-1} [v_{n-1} - b_{n-1,n} \tilde{\eta}_n]$. Evaluate the cumulative standard-normal at this upper bound and call it p_{n-1} . Then use the first two draws to calculate the $(n-2)^{\text{th}}$ upper bound and calculate p_{n-2} analogously, and so on through all n observations. Formally, this recursive process for calculating the upper bounds is:

$$\left. \begin{aligned} \eta_n &< b_{n,n}^{-1} v_n \equiv v_n \\ \eta_{n-1} &< b_{n-1,n-1}^{-1} [v_{n-1} - b_{n-1,n} \tilde{\eta}_n] \equiv v_{n-1} \\ \eta_{n-2} &< b_{n-2,n-1}^{-1} [v_{n-2} - b_{n-2,n-1} \tilde{\eta}_{n-1} - b_{n-2,n} \tilde{\eta}_n] \equiv v_{n-2} \\ &\vdots \end{aligned} \right\} \Rightarrow \eta_j < b_{j,j}^{-1} \left[v_j - \sum_{i=j+1}^n b_{j,i} \tilde{\eta}_i \right] \equiv v_j \quad (21).$$

The probability of observing a given sample of ones and zeros can now be found by evaluating the *univariate* cumulative-normal distribution function at each of these bounds, p_i , and then multiplying

those probabilities: $\prod_{j=1}^n p_j = \prod_{j=1}^n \Phi(v_j)$. Repeating the entire process R times and averaging gives the

RIS estimate for the joint probability, i.e., the simulated likelihood, as this mean:

$$\hat{l} = (1/R) \sum_{r=1}^R \left[\prod_{j=1}^n \Phi(v_{j,r}) \right] \quad (22).$$

One can then maximize this simulated likelihood by any standard optimization routine to estimate parameters and apply the standard ML estimator for the variance-covariance (i.e., $-[H(\hat{l})]^{-1}$).

IV. Monte Carlo Analyses of Standard-Probit vs. Bayesian-Spatial-Probit Estimation

We explore the small-sample properties of standard ML-probit and Bayesian MCMC estimators for the spatial-lag probit model using a data-generation process (DGP) that closely follows Beron & Vijverberg's (2004) Monte Carlo exploration of the RIS estimator:

$$\mathbf{y}^* = (\mathbf{I}_n - \rho \mathbf{W})^{-1} (\mathbf{x}' \boldsymbol{\beta} + \boldsymbol{\varepsilon}), \text{ where } \mathbf{x} = (\mathbf{I}_n - \theta \mathbf{W})^{-1} \mathbf{z} \text{ and } \mathbf{z}, \boldsymbol{\varepsilon} \sim N(0,1) \quad (23).$$

We apply (3) to generate \mathbf{y} from these \mathbf{y}^* . Note that \mathbf{x} and \mathbf{y}^* exhibit the same pattern of spatial interdependence, \mathbf{W} , but with strengths θ vs. ρ . For \mathbf{W} , we use a row-standardized binary-contiguity matrix for the 48 contiguous U.S. states. We set ρ to 0.5 and β to 1.0, and consider sample sizes $n = \{48, 144\}$ ²⁸ and θ values 0.0 and 0.5, giving four experiments total. Table 1 reports the results for

²⁸ To create the weights matrix for the larger sample size we took the Kronecker product of the original 48×48 weights matrix with a 3×3 identity matrix. This could reflect, e.g., three observations of outcomes in each of the 48 states.

1,000 trials using the standard-probit ML estimator with spatial lags $\mathbf{W}y$ or $\mathbf{W}y^*$, and the Bayesian MCMC spatial-probit estimator.²⁹ Standard ML-probit with a $\mathbf{W}y^*$ regressor is not practicable because y^* is unobserved, but those results provide us valuable comparison. Only the simultaneity of the spatial lag biases those estimates, whereas the incorrect $\mathbf{W}y$ spatial-lag used in current standard-practice incurs both simultaneity, with its likely inflation-bias, and measurement error (of $\mathbf{W}y$ vs. $\mathbf{W}y^*$), with its attenuation bias. Beron & Vijverberg (2004) thoroughly evaluate the RIS estimator's properties using a similar DGP, so we refer to those results for RIS rather than reanalyze it here.

Table 1: Simulation Results

	ML with $\mathbf{W}y$		ML with $\mathbf{W}y^*$		Bayesian MCMC	
	β	ρ	β	ρ	β	ρ
Experiment #1: n=48, $\theta=0.0$						
Mean Coefficient Estimate	1.02	0.32	1.13	0.74	1.23	0.30
Root Mean-Squared Error	0.33	0.71	0.43	0.43	0.36	0.26
Actual Std Dev of Estimates	0.33	0.69	0.41	0.36	0.28	0.16
Mean of Reported Std Err	0.30	0.41	0.35	0.30	0.42	0.21
Experiment #2: n=48, $\theta=0.5$						
Mean Coefficient Estimate	1.22	0.35	1.13	0.69	1.21	0.28
Root Mean-Squared Error	0.60	0.77	0.62	0.38	0.32	0.26
Actual Std Dev of Estimates	0.56	0.76	0.61	0.33	0.24	0.14
Mean of Reported Std Err	0.36	0.46	0.42	0.29	0.39	0.20
Experiment #3: n=144, $\theta=0.0$						
Mean Coefficient Estimate	0.94	0.42	1.01	0.68	1.14	0.34
Root Mean-Squared Error	0.18	0.28	0.19	0.24	0.21	0.19
Actual Std Dev of Estimates	0.17	0.27	0.19	0.16	0.15	0.10
Mean of Reported Std Err	0.16	0.22	0.18	0.15	0.22	0.12
Experiment #4: n=144, $\theta=0.5$						
Mean Coefficient Estimate	1.08	0.48	0.97	0.64	1.13	0.32
Root Mean-Squared Error	0.21	0.29	0.21	0.21	0.19	0.20
Actual Std Dev of Estimates	0.19	0.29	0.21	0.16	0.14	0.09
Mean of Reported Std Err	0.18	0.23	0.20	0.15	0.21	0.12

Perhaps the most surprising result from our experiments is the relatively poor bias performance of the Bayesian MCMC estimator.³⁰ The common-practice ML-probit with spatial lag $\mathbf{W}y$ actually weakly dominates on bias criteria (of $\hat{\rho}$ and, less strikingly, $\hat{\beta}$), although its much lesser efficiency

²⁹ We use LeSage's (1999) MatLab code, with q set to the default probit value determinately (see note 24), and a burn-in of 1000 trials, retaining the next 1000 for our simulation sample.

³⁰ We recognize that this evaluates a Bayesian estimator by frequentist standards, but we think those standards worth considering nonetheless. We also have not yet explored whether the problem is one of the estimator, our implementation

makes the MCMC estimator conversely near-dominant by root mean-squared error. We know that the common-practice estimator suffers two biases, but those seem fortuitously to offset somewhat in these experiments. The first is the simultaneity bias that also plagues the ML estimator with the true spatial-lag, $\mathbf{W}\mathbf{y}^*$. In that latter case, this is the only source of bias, and, indeed, those columns show strong inflation of $\hat{\rho}$. The second bias of the common-practice standard ML-probit is an attenuation bias in $\hat{\rho}$ due to measurement error in proxy spatial-lag, $\mathbf{W}\mathbf{y}$, compared to true spatial-lag, $\mathbf{W}\mathbf{y}^*$. The simultaneity inflation-bias increases with ρ , but the impact of the attenuation bias instead decreases with sample size (for this particular \mathbf{W} at least). Therefore, when ρ and n are small, measurement-error attenuation dominates, leaving a net-negative bias. When ρ and n are large, the simultaneity inflation-bias dominates, and net bias is positive. This implies that at some ρ and n between (somewhere near the conditions of our fourth experiment, apparently) the biases cancel.³¹

Notice also that where \mathbf{x} correlates spatially (and, as here, in a pattern like that of \mathbf{y}), over/under-estimates of ρ generally associate with under/over-estimates of β , as we had found for the linear-regression case, although not as consistently or proportionately here. Thus, in binary-outcomes too, omission or inadequate modeling of spatial interdependence will tend to bias conclusions toward non-spatial (unit-level or exogenous-external) explanations, and *vice versa*.

In our smaller sample, the standard errors from the ML with proxy spatial-lag are overconfident about the estimator's precision, with standard-error estimates for $\hat{\rho}$ underestimated in both cases by 40%. The standard ML estimator with the true spatial-lag also overstates confidence, though by only about half as much. The Bayesian MCMC estimator, contrarily, overestimates uncertainty by about a third in the smaller-sample cases. In all cases, misestimation of uncertainty diminishes with sample size. Biases in the standard-error estimate for $\hat{\beta}$, meanwhile, are only noticeable where the x 's are

³¹ Further preliminary experiments varying ρ and n , using fewer trials for speed, so far confirm these intuitions.

spatial interdependent. Under conditions like those of our second experiment, Beron and Vijverberg (2004, Tables 8.3 & 8.4) report that RIS overestimates β by 10% and underestimates ρ by 18%, which compares favorably to the same biases for the standard ML estimator with proxy spatial-lag (22% and 30%) and at least as well to those for the Bayesian MCMC in this case (21% and 44%).

V. Calculating and Presenting Estimated Spatial Effects with Certainty Estimates

Properly estimating *parameters* such as *coefficients* and their certainties is obviously essential to valid inference, but our ultimate aims usually are to estimate, draw inferences regarding, interpret and present *effects* (ideally: causal ones), i.e., changes in the expectations of *outcomes* associated with (ideally: caused by) changes in explanatory factors or other counterfactual shocks. Ultimately, we estimate *coefficients* like ρ and β for the purpose of estimating *effects* like $\frac{\partial y_i^*}{\partial x_i}$ or $\frac{\Delta y_i^*}{\Delta x_i}$, i.e., the effects of a marginal or discrete change in some explanatory factor in unit i , x_i , on the latent-variable outcome in i , y_i^* , or, better, the effects of x_i on the probability of i 's choice or outcome, $\frac{\partial p(y_i=1)}{\partial x_i}$ or $\frac{\Delta p(y_i=1)}{\Delta x_i}$. Given interdependence, even these sorts of “within-unit” counterfactuals involve feedback from i through other units j back to i . In fact, in diffusion, interdependence, or spatial- or network- interaction contexts (roughly synonyms) our interests usually extend centrally to the cross-unit feedback effects, such as $\frac{\partial y_j^*}{\partial x_i}$, $\frac{\Delta y_j^*}{\Delta x_i}$, $\frac{\partial p(y_j=1)}{\partial x_i}$, or $\frac{\Delta p(y_j=1)}{\Delta x_i}$. Either within or across units, we could also wish to consider some generic shocks to y_i^* , the linear “propensity” toward outcome $y_i=1$, rather than shock to some x_i . For these purposes we expanding the latent model to include some unspecified unit-specific factor, i.e., “unit effects” (fixed not random), ω_i , to $\mathbf{y}^* = \rho \mathbf{W} \mathbf{y}^* + \mathbf{X} \beta + \boldsymbol{\omega} + \boldsymbol{\varepsilon}$. Finally, in interdependent binary-outcome contexts, we are likely to want estimates of the effects on the probability of some unit(s) i 's choices/outcomes of counterfactual shocks to choices/outcomes of other unit(s) j . For instance, the effect on the probability Michigan enacts some policy of Illinois

and/or Ohio enacting it. We denote this sort of counterfactual effect as $\frac{\partial y_i}{\partial y_j}$ or $\frac{\Delta y_i}{\Delta y_j}$. In contexts of spatially interdependent binary outcomes, none of these substantive *effects* is simple to estimate; indeed, the difficulty of calculation increases directly with the centrality of their substantive interest.

To begin, we remind and emphasize that only in purely linear and additively separable models, like the canonical regression, $y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \varepsilon$, are *coefficients* and *effects* (of changes in x on y) identical. Even in models only implicitly nonlinear-additive, like spatial-autoregressive linear-regression, *effects* involve (often nonlinear) combinations of coefficients and variables, *via* the spatial-feedback multipliers in that case. Thus, even if we were content to confine our interpretation and presentation to the latent-variable arguments, y^* , to the probabilities of actual interest, $\hat{\mathbf{p}}$, we could not read *effects* directly from the usual table of *coefficients*. Instead, calculation, interpretation, and presentation of estimated effects on latent variables and their certainties would ensue as in the spatial linear-regression that we have discussed extensively elsewhere. To review:

$$\mathbf{y}^* = \rho \mathbf{W} \mathbf{y} + \mathbf{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon} = (\mathbf{I}_n - \rho \mathbf{W})^{-1} (\mathbf{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon})$$

$$= \begin{bmatrix} 1 & -\rho w_{1,2} & \cdots & \cdots & -\rho w_{1,n} \\ -\rho w_{2,1} & 1 & & & \vdots \\ \vdots & & \ddots & & \vdots \\ \vdots & & & 1 & -\rho w_{n-1,n} \\ -\rho w_{n,1} & \cdots & \cdots & -\rho w_{n,n-1} & 1 \end{bmatrix}^{-1} (\mathbf{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon}) \equiv \mathbf{S} (\mathbf{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon}) \quad (24).$$

Thus, denoting the i^{th} column of \mathbf{S} as \mathbf{s}_i and their estimates as $\hat{\mathbf{S}}$ and $\hat{\mathbf{s}}_i$, the estimated effect of explanatory variable k in unit i , $\Delta x_{i,k}$, on the outcomes in all units, i and all j , is $\frac{\Delta \hat{\mathbf{S}} \mathbf{X} \hat{\boldsymbol{\beta}}}{\Delta x_{i,k}}$ which is simply, $\hat{\mathbf{s}}_i \hat{\boldsymbol{\beta}}_k$. The standard-error calculation, using the delta method approximation, is

$$\hat{V}(\hat{\mathbf{s}}_i \hat{\boldsymbol{\beta}}_k) \approx \left[\frac{\partial \hat{\mathbf{s}}_i \hat{\boldsymbol{\beta}}_k}{\partial \hat{\boldsymbol{\theta}}} \right] \hat{V}(\hat{\boldsymbol{\theta}}) \left[\frac{\partial \hat{\mathbf{s}}_i \hat{\boldsymbol{\beta}}_k}{\partial \hat{\boldsymbol{\theta}}} \right]', \text{ where } \hat{\boldsymbol{\theta}} = \begin{bmatrix} \hat{\rho} \\ \hat{\boldsymbol{\beta}}_k \end{bmatrix} \text{ and } \left[\frac{\partial \hat{\mathbf{s}}_i \hat{\boldsymbol{\beta}}_k}{\partial \hat{\boldsymbol{\theta}}} \right] = \begin{bmatrix} \frac{\partial \hat{\mathbf{s}}_i \hat{\boldsymbol{\beta}}_k}{\partial \hat{\rho}} & \hat{\mathbf{s}}_i \end{bmatrix} \quad (25).$$

The vector $\frac{\partial \hat{\beta}_k}{\partial \hat{\rho}}$ is the i^{th} column of $\beta_k \frac{\partial \hat{\mathbf{S}}}{\partial \hat{\rho}}$. Since \mathbf{S} is an inverse matrix, the derivative in equation

$$(25) \text{ is } \frac{\partial \hat{\mathbf{S}}}{\partial \hat{\rho}} = -\hat{\mathbf{S}} \frac{\partial \hat{\mathbf{S}}^{-1}}{\partial \hat{\rho}} \hat{\mathbf{S}} = -\hat{\mathbf{S}} \frac{\partial (\mathbf{I} - \hat{\rho} \mathbf{W})}{\partial \hat{\rho}} \hat{\mathbf{S}} = -\hat{\mathbf{S}} (-\mathbf{W}) \hat{\mathbf{S}} = \hat{\mathbf{S}} \mathbf{W} \hat{\mathbf{S}}.$$

Elsewhere, we showed how to use these and related expressions to generate grids, tables, or maps of responses across units to various counterfactuals, along with appropriate indicators of the estimated certainties of these estimated spatial effects. We also showed in the spatiotemporal context how to estimate and graph spatiotemporal response-paths and estimate and tabulate or array in grids long-run-steady-state spatiotemporal responses to counterfactuals, along with certainty estimates thereof.

If we confine our attention to the latent variable, \mathbf{y}^* , all of these techniques could apply in the spatial-probit context exactly as previously described, but, for most purposes, interpretation in terms of latent \mathbf{y}^* is unsatisfactory. Furthermore, several issues regarding the application of delta-method asymptotic linear-approximation increasingly trouble us, the intrinsic appeal of analytic solutions notwithstanding. First, deriving from a linearization, the certainty estimates only approximate validly in some proximity of the estimated *nonlinear* expression, and we do not know in general how small a range. Being asymptotic, they only approximate validly for large samples, and we do not know in general how large, and they are in any event an approximation. Finally, using the approximately estimated standard errors to generate confidence intervals and hypothesis tests in the usual manners assumes (multivariate) normality of the parameter estimates. In maximum-likelihood contexts, this is not especially problematic since all ML estimates are at least asymptotically normal, though sample-size concerns may arise, perhaps especially regarding estimates involving $\hat{\rho}$, which is exactly where the spatial complications tend to arise. Given all this, we increasingly suspect that the asymptotic linear-approximations we have been recommending may have been larger than need be even in the linear-regression context. For those spatial linear-regression contexts, simple simulation strategies—

i.e., sampling coefficient estimates from multivariate normal with the estimated means and variance-covariance, calculating the quantities of interest from those draws, and then generating the desired indicators of certainty from the resulting sample—may be more effective.

Even greater concerns arise in the spatial-probit context because the nonlinearity of the estimates of interest is more severe and asymptotic normality may be more distant. In fact, the (kernel of the) posterior joint-distribution of the parameters is not normal (as seen in (10), due to the $|\mathbf{I}_n - \rho\mathbf{W}|$ term), and, of the posterior conditional-distributions, only that of $\boldsymbol{\beta}$ is exactly normal. Thus, we suggest using the same MCMC (or RIS) processes that yielded the parameter estimates and their certainty estimates to estimate by simulation the quantities on interest and their certainties. To elaborate, recall that, after sufficient burn-in, LeSage’s Metropolis/Hastings-within-Gibbs sampler generates draws from the poster joint-distribution of the parameters. The parameter estimates are the sample-means of these draws, and certainty estimates for those parameter-estimates are variances or percentile-ranges of those draws. Since one property of the Gibbs sampler is that it converges to the correct joint-posterior of the parameters, we could simply calculate any quantity of interest for each element of the (post-burn-in) sample vector of parameters, supplying whatever counterfactual values of interest for whatever variables enter the expression of interest. The RIS-simulated likelihoods would support the same procedure, but a serious complication would yet remain in either case.

To begin, consider our interests in levels or changes of \hat{p}_i and \hat{p}_j ’s or, most generally, $\hat{\mathbf{p}}$, the vector of probabilities of 1’s in units i and j induced by hypothetical levels or changes in some $x_{i,k}$ or $x_{j,k}$, or, most generally, \mathbf{X} . For instance, using (4), we could calculate the effects of some change in \mathbf{X} on the estimated probability of an outcome of 1 in unit i as:

$$\frac{\Delta[p(y_i = 1)]}{\Delta\mathbf{X}} = p\left(u_i < \frac{[(\mathbf{I} - \rho\mathbf{W})^{-1}\mathbf{X}_1\boldsymbol{\beta}]_i}{\sigma_i}\right) - p\left(u_i < \frac{[(\mathbf{I} - \rho\mathbf{W})^{-1}\mathbf{X}_0\boldsymbol{\beta}]_i}{\sigma_i}\right) \quad (26).$$

where $\Delta\mathbf{X} = \mathbf{X}_1 - \mathbf{X}_0$ is the hypothetical change being considered in some x or x 's in some unit(s). Notice that to calculate the effect even of a change in one x in one unit i on the outcome in just that i , the researcher must specify not only the from/to levels of that change and the levels of all the \mathbf{x}_i , as in standard probit, but also all the levels of all the \mathbf{x}_j in all the other units. Intuitively, this is because not only do all the \mathbf{x}_i affect where we are on the probit sigmoid curve, as usual, but all the y_j^* also affect that positioning *via* spatial feedbacks, and those in turn depend on all \mathbf{x}_j (and all the other y_{-j}^* , including y_i^* , and so on). These expressions and procedures hold for any i and $\Delta\mathbf{X}$, so the cross-unit effects on some j of changes in some i are calculated by the same formula, applying the desired $\Delta\mathbf{X}$ and changing the subscripts to refer to j^{th} elements. This seems feasible, although the need to specify all of $\Delta\mathbf{X}$ for any counterfactual may be a bit daunting, but a far larger challenge is still looming.

Just as in the estimation problem, the $p(u_i < [(\mathbf{I} - \hat{\rho}\mathbf{W})^{-1}\mathbf{X}\hat{\boldsymbol{\beta}}]_i / \hat{\sigma}_i)$ of interest here emerge from a multivariate cumulative-normal with means $\mathbf{0}$ and variance-covariance $[(\mathbf{I} - \rho\mathbf{W})'(\mathbf{I} - \rho\mathbf{W})]^{-1}$. In the case of estimation, we sought to maximize a likelihood conditional on the data, i.e., \mathbf{y} and \mathbf{X} , which implied that we needed to evaluate one n -dimensional cumulative normal rather than multiply n unidimensional cumulative normals. To understand exactly how the same issue arises in estimating our counterfactual effects, consider the following spatial-probit model, simplified to a bivariate case:

$$\begin{aligned} y_1^* &= \rho w_{12} y_2^* + \beta_1 x_1 + \eta_1 + \varepsilon_1 \\ y_2^* &= \rho w_{21} y_1^* + \beta_2 x_2 + \eta_2 + \varepsilon_2 \end{aligned} \quad (27),$$

with η_i a fixed effect specific to y_i^* and $\varepsilon_i \sim N(0,1)$. The reduced form of the model is:

$$\begin{aligned} y_1^* &= \frac{\beta_1}{1-\rho^2 w_{12} w_{21}} x_1 + \frac{\rho w_{12} \beta_2}{1-\rho^2 w_{12} w_{21}} x_2 + \frac{1}{1-\rho^2 w_{12} w_{21}} \eta_1 + \frac{\rho w_{12}}{1-\rho^2 w_{12} w_{21}} \eta_2 + \frac{1}{1-\rho^2 w_{12} w_{21}} \varepsilon_1 + \frac{\rho w_{12}}{1-\rho^2 w_{12} w_{21}} \varepsilon_2 \\ y_2^* &= \frac{\rho w_{21} \beta_1}{1-\rho^2 w_{12} w_{21}} x_1 + \frac{\beta_2}{1-\rho^2 w_{12} w_{21}} x_2 + \frac{\rho w_{21}}{1-\rho^2 w_{12} w_{21}} \eta_1 + \frac{1}{1-\rho^2 w_{12} w_{21}} \eta_2 + \frac{\rho w_{21}}{1-\rho^2 w_{12} w_{21}} \varepsilon_1 + \frac{1}{1-\rho^2 w_{12} w_{21}} \varepsilon_2 \end{aligned} \quad (28).$$

Latent y_i^* still links to the observed binary variable y_i through measurement equation (3), implying:

$$y_i = \begin{cases} 1 & \text{if } \varepsilon_i + \rho w_{ij} \varepsilon_j < \beta_i x_i + \rho w_{ij} \beta_j x_j + \eta_i + \rho w_{ij} \eta_j \\ 0 & \text{if } \varepsilon_i + \rho w_{ij} \varepsilon_j \geq \beta_i x_i + \rho w_{ij} \beta_j x_j + \eta_i + \rho w_{ij} \eta_j \end{cases} \quad (29)$$

The joint probability of any y_1 and y_2 is the product of a marginal and conditional probability; e.g.:

$$\Pr(y_1 = 1 \cap y_2 = 1) = \Pr(y_2 = 1 | y_1 = 1) \times \Pr(y_1 = 1) \quad (30)$$

For estimation purposes, given sample observations on y_1 and y_2 , we apply the appropriate version of (30)'s right-hand side to calculate the joint likelihood for the pair of observations. One sees this product of one marginal and $n-1$ conditional distributions directly in the RIS estimator's equation (19), for example. If we wanted to calculate the marginal probability for y_1 , $\Pr(y_1 = 1)$, i.e., the probability $y_1 = 1$ unconditional on ε_2 , i.e., unconditional on the other unit, i.e., unconditional on y_2

or the $\Pr(y_2 = 1)$, we would integrate over ε_2 . Then, because $\rho w_{21} \int_{-\infty}^{\infty} \varepsilon_2 f(\varepsilon_2) d\varepsilon_2 = 0$, the $\rho w_{21} \varepsilon_2$ term of (29) drops from the calculation, which means the simple univariate cumulative normal could be evaluated at the right-hand-side value to obtain $\Pr(y_1 = 1)$. That is, the marginal probability for y_1 depends on x_2 and η_2 (and x_1 and η_1 , of course), but not on ε_2 , the disturbance term from y_2^* .

However, the essence of interdependence would suggest that we are not particularly interested in these marginal probabilities, substantively. We want to consider counterfactual shocks to \mathbf{X} or $\boldsymbol{\eta}$, including the feedback represented in $\mathbf{W}\mathbf{y}^*$, which means conditional on ε_2 . Calculating conditional probabilities like $\Pr(y_1 = 1 | y_2 = 1)$ is more complicated because this probability depends on the disturbance term from y_2^* . Since we are conditioning on $y_2 = 1$ (in this example), the possible error term from y_2^* , call it $\tilde{\varepsilon}_2$, is a random variable that comes from a truncated normal distribution with support over the range $[-\infty, \beta_1 x_1 + \rho w_{12} \beta_2 x_2 + \eta_1 + \rho w_{12} \eta_2]$. Since these distribution are truncated at

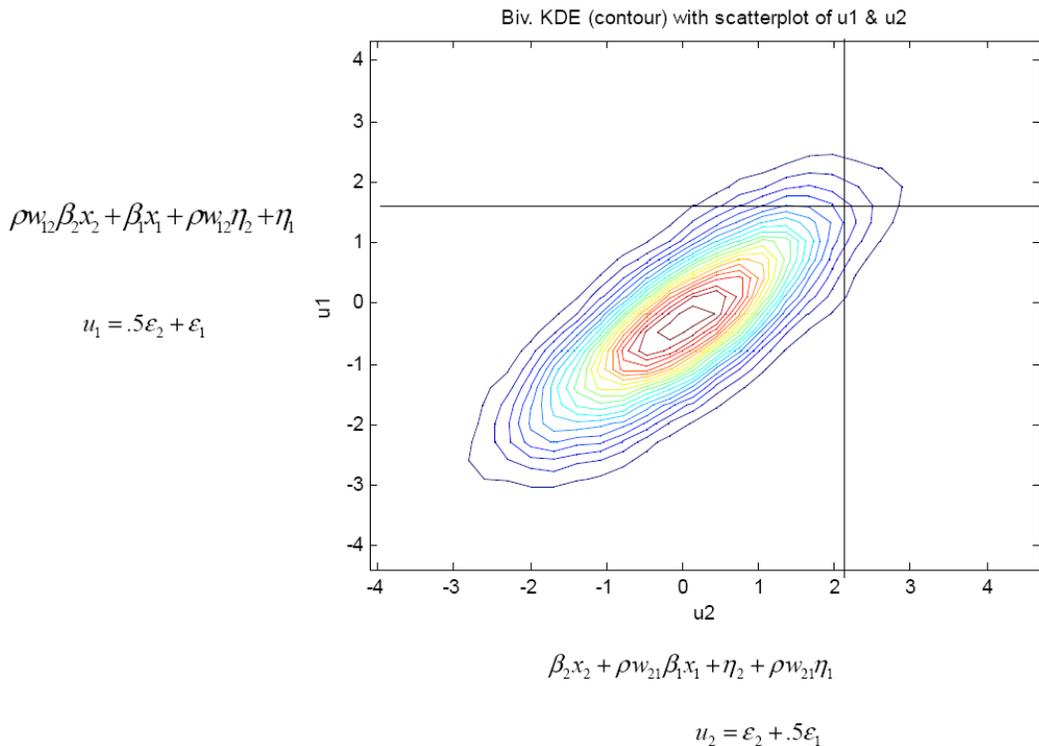
the cutpoints for the conditional effects, the $\rho w_{21}\varepsilon_2$ term of (29) does *not* drop from the calculation, and we must compute the n-dimensional cumulative normal, just as in the original estimation stage.

More specifically, the marginal probabilities are

$$\begin{aligned} \Pr(y_1 = 1 | \mathbf{x}, \boldsymbol{\eta}) &= \Pr(\varepsilon_1 < \beta_1 x_1 + \rho w_{12} \beta_2 x_2 + \eta_1 + \rho w_{12} \eta_2) \\ &= \Phi[\beta_1 x_1 + \rho w_{12} \beta_2 x_2 + \eta_1 + \rho w_{12} \eta_2] \end{aligned} \quad (31)$$

$$\begin{aligned} \Pr(y_2 = 1 | y_1 = 1; \mathbf{x}, \boldsymbol{\eta}) &= \Pr(\rho w_{21} \tilde{\varepsilon}_1 + \varepsilon_2 < \beta_2 x_2 + \rho w_{21} \beta_1 x_1 + \eta_2 + \rho w_{21} \eta_1) \\ &= \Phi[\beta_2 x_2 + \rho w_{21} \beta_1 x_1 + \eta_2 + \rho w_{21} \eta_1 - \rho w_{21} \tilde{\varepsilon}_1] \end{aligned} \quad (32).$$

Both of these cumulative distribution-functions are of the unidimensional standard-normal but, in the case of (32), because the left-hand-side term in line one involving ε has been transformed by partial differencing (and multiplication by the denominator in (28), which would be retained on the right-hand side also). To get an unbiased estimate of the conditional probability in (32), we can take a draw from the truncated normal distribution for values of $\tilde{\varepsilon}_1$ (see below). Taking R draws, and averaging the probabilities, enhances the efficiency of this maximum simulated likelihood estimator.



To reiterate, to calculate counterfactual effects of shocks to some units on probabilities of outcomes in some units, we could consider marginal and conditional probabilities. The marginal probabilities for \mathbf{y}_i do not depend on \mathbf{y}_j , though they do depend on the full matrix \mathbf{X} and vector $\boldsymbol{\eta}$. The conditional probabilities for \mathbf{y}_i , which have more substantive meaning, do depend on \mathbf{y}_j , as well as the full matrix \mathbf{X} and vector $\boldsymbol{\eta}$, but are possible, though more difficult to calculate. In short, estimating effects in terms of probabilities of outcomes, i.e., in terms of the substantive quantity of interest, is as computationally burdensome as obtaining the estimates, for exactly the same reason. This effect-estimate procedure must be repeated many times to estimate their variance-covariance.

In principle, then, we can calculate the Δp_i responses in all units, $\Delta \mathbf{p}$, for any hypothetical change, $\Delta \mathbf{X}$, by this formula:

$$\frac{\Delta \mathbf{p}}{\Delta \mathbf{X}} = \boldsymbol{\Phi}_n \left(\left[(\mathbf{I} - \rho \mathbf{W})^{-1} \mathbf{X}_1 \boldsymbol{\beta} \right] \odot \left[\left\{ \sigma_i^{-1} \right\} \right] \right) - \boldsymbol{\Phi}_n \left(\left[(\mathbf{I} - \rho \mathbf{W})^{-1} \mathbf{X}_0 \boldsymbol{\beta} \right] \odot \left[\left\{ \sigma_i^{-1} \right\} \right] \right) \quad (33),$$

with $\Delta \mathbf{p}$ the $n \times 1$ vector of $\Delta [p(y_i = 1)]$ across all i ; $\boldsymbol{\Phi}_n(\bullet)$ the cumulative-normal distributions, evaluated element-by-element at the values of its $n \times 1$ vector argument, from the n -variate normal distribution with means zero and variance-covariance $[(\mathbf{I} - \rho \mathbf{W})'(\mathbf{I} - \rho \mathbf{W})]^{-1}$; $\left[\left\{ \sigma_i^{-1} \right\} \right]$ the $n \times 1$ vector of the previously defined scalars σ_i^{-1} ; and \odot indicating element-by-element multiplication (i.e., Hadamard product). In principle, for given $\hat{\rho}$, these integrals could be calculated numerically using RIS or Gibbs sampling techniques, and certainty estimates for these effect estimates could then be obtained by repeating the process for many draws. However, calculating effects this way would take c times as long as the estimation procedure, with c the number of effect-estimates from which the estimated variance of the effect-estimate derives; computational intensity would be prohibitive.³²

³² A simpler expedient may exist to evade integration of the n -dimensional multivariate-normal by drawing coefficients

VI. Illustration: Diffusion of Legislative Term-Limits among the American States

As illustration, we draw on the policy-diffusion and states-as-laboratories literatures in the study of U.S. politics (e.g., Volden 2006, Morehouse & Jewell 2004). Specifically, we consider a spatial-lag model of term-limit adoption in the states. The spatial lag allows us to consider whether states learn from or are otherwise influenced by their neighbors. Many studies examine the effects of term limits on the composition and functioning of legislatures or on the behavior of individual legislators (e.g., Carey et al. 1998, Cain & Levin 1999), but why states might adopt limits in the first place has received much less scholarly attention, adding to the interest of the example. The dependent variable in our analysis indicates (0,1) whether a state has adopted term limits. From 1990 to 2000, 21 states adopted term limits.³³ The principal determinant of term limit adoption is whether a state allows ballot initiatives or popular referenda (I&R) to consider state-level statutes and/or constitutional amendments (i.e., direct democracy). Simple reasoning likely underlies this strong relationship.

State legislators, particularly career politicians, are less likely than the public to want term limits, and direct democracy allows the electorate to bypass the legislature (Cain & Levin 1999). Indeed, only one state that does not allow such direct democratic processes, Louisiana, has term limits. Table 2 lists the 27 states (of the 48 contiguous) that either have term-limits or allow some form of direct democracy. The relationship is extremely strong: 21 of the 22 states without I&R do not have term

from the multivariate posterior or simulated-likelihood of σ , ρ , and β . (If one wishes to include estimated inherent-uncertainty as well as estimation-uncertainty in these counterfactuals, then one should also draw ϵ from its independent-normal distributions, adding it to the $\mathbf{X}\hat{\beta}$ in the next term.) Calculate \hat{y}_1^* and \hat{y}_0^* using $(\mathbf{I}_n - \rho\mathbf{W})^{-1} \mathbf{X}\hat{\beta}$ for some fixed \mathbf{X}_1 and \mathbf{X}_0 , then simply apply (3) to convert those to vectors of ones and zeros, \hat{y}_1 and \hat{y}_0 . For a large number of draws, the averages of \hat{y}_1 and \hat{y}_0 will be $\hat{\mathbf{p}}_1$ and $\hat{\mathbf{p}}_0$, and $\hat{\mathbf{p}}_1 - \hat{\mathbf{p}}_0$ will be the desired vector of estimated effects, and the variance-covariance of those differences will be the variance-covariance of those estimated effects. We are still working through some problems with this conjecture, however.

³³ In six of these states, term limits have either been overturned by state supreme courts or repealed by state legislatures. We code the dependent variable as 1 in these six cases.

limits, and 20 of the 26 states with I&R do have them.³⁴ The six states with I&R but without term limits are Illinois, Kentucky, Maryland, Mississippi, New Mexico, and North Dakota.

Our empirical models include two other explanatory variables. The first indicates (0,1) whether the state voted for Clinton in the 1992 presidential election. Some argue that Democrats, because of their relatively positive view of government and related support for state intervention, are more accepting of political careerism and therefore more likely to oppose term limits.³⁵ On the other hand, the populist tendencies of some Democrats may lead them to support term limits as a way to promote citizen participation in government. The second is the average state-level tax effort (state revenue as percent of state GDP) during the 1980s. High state taxes may indicate public support for political centralization and a strong professionalized legislature, or, alternatively, high taxes may provide impetus for an overburdened electorate to “throw the bums out” using term limits.

Table 3 reports estimates of probit models with spatial-lag regressors (or, in one column, spatial error-dependence) by standard maximum-likelihood (ML) methods that erroneously assume spatial lags exogenous, by Bayesian MCMC methods but maintaining the same erroneous assumption, and by true spatial-lag probit (or spatial-error probit) using the Bayesian MCMC and the frequentist RIS methods described in Section III. We use a standardized binary contiguity-weights matrix, \mathbf{W} , which codes $w_{ij} = (1,0)$ for whether states i and j border and then row-standardizes the resulting matrix by dividing each element by its row’s sum.³⁶ This gives $(\mathbf{W}\mathbf{y})_i$ as the unweighted average of the outcome in i ’s bordering states—i.e., the share of bordering states that have term limits—or

³⁴ We can easily reject the null hypothesis that these two variables are independent ($\chi^2_{(1)} = 25.367$, p-value=.000), and Kendall’s τ_b , a (-1...+1) correlation-like measure of association, is .727 with asymptotic standard error of .091.

³⁵ There is individual-level evidence for a relationship between Republican partisanship and support for term limits in several states (see Cain and Levin 1999 for a discussion).

³⁶ Row-standardization is standard in spatial econometrics, but it is not necessarily substantively neutral (see, e.g., Pluemper & Neumayer 2008).

$(\mathbf{W}\mathbf{y}^*)_i$ as the unweighted average propensity of i 's neighbors to adopt term limits.

[Table 3 Here]

The first two columns report models estimated (wrongly) assuming the spatial lags exogenous. The first-column model applies standard-probit ML techniques. The parentheses contain the standard estimated standard errors, with the hypothesis tests assuming the test-statistics asymptotic-normally distributed. The next two columns' models are estimated using MCMC methods with diffuse zero-mean priors, including an uninformative uniform(-1,1) prior on ρ . The reported coefficient-estimates are means of posterior distributions using 10,000 cycles of the sampler after a 1000-cycle burn-in. The parentheses report sample standard-deviations of the posterior distributions, and p -values also emerge directly from the posterior (without calculating or assuming anything about test statistics).

The results in columns one and two are similar, which is not surprising given our use of diffuse priors. The results in columns two and three are more noticeably different. This is because the probit-MCMC estimator used in column two, as with probit-ML, incorrectly treats the spatial lag as exogenous (i.e., as any other right-hand-side variable). Therefore the likelihood is misspecified, so the sampler draws from the wrong posterior distribution for the spatial coefficient $\hat{\rho}$. As we have seen, these specification errors seriously compromise inferences from either of these models about the strength and importance of spatial interdependence. The result here seems an overestimation by more than 300% of interdependence-strength and a more than three-fold overestimation of the uncertainty, judging the second relative to the third column. The posterior-distribution over the direct democracy (I&R) parameter is also notably influenced by this misspecification, with the marginal distribution in column two having a much smaller mean and variance than the one in column three.

Column three reports the Bayesian-Gibbs spatial-probit estimates. The draws for $\hat{\rho}$ are taken from the correct (non-standard) posterior distribution using Metropolis-Hastings. About 24% of the

10,000 spatial AR coefficients sampled from the posterior distribution were less than zero. Given our diffuse priors, forty-eight observations cannot produce very sharp posterior beliefs about the degree (or even the sign) of spatial interdependence. Direct democracy remains an important determinant using this estimator. In fact, the mean of the posterior distribution over the coefficient on I&R is larger than any of the other posterior means or I&R coefficient estimates. The spatial-error model produces similar results with respect to the state-level variables, but it also suggests much weaker interdependence. The problem with the spatial-error model, though, is that the specification simply lacks any substantive sense. One cannot easily imagine a reason why the stochastic component of y^* should exhibit spatial interdependence while the systematic component does not.

Finally, the RIS estimates produce a similar pattern of coefficient-estimates, with $\hat{\rho}$ between that of the two Gibbs-estimated models (columns 2 and 3); however, the RIS estimator finds $\hat{\rho}$ to be rather clearly statistically significant. That the RIS and Bayesian-Gibbs results give such different answers regarding the statistically discernable importance of spatial interdependence is troubling. Both these methods use correctly specified likelihoods built upon substantively plausible models. The only difference between the two methods is that the posterior distribution is non-standard in the Bayesian case, whereas the sampling distribution for the RIS estimator is multivariate normal. The updating from diffuse prior to non-standard posterior seems to be less efficient with the Bayesian estimator than the similar “updating” from diffuse (implicit) prior to standard “posterior” (sampling) distribution with the RIS estimator. If we calculate the likelihood using the coefficient estimates in column 3 we get -17.943, which is much less than the value of the likelihood at the RIS parameter estimates (-16.261). This seems to suggest that the data are having less influence on the Bayesian posterior distribution than on the frequentist sampling distribution. This interpretation of the results in Table 3 is also consistent with our and Beron and Vijverberg’s (2004) Monte Carlo results. In the

comparable simulations, we find the Bayesian-Gibbs estimator underestimates the true value of ρ by 44% on average, whereas Beron and Vijverberg find that the RIS estimator underestimates ρ by only 18% on average. Given the superior likelihood value and the available Monte Carlo evidence, we are more confident in the RIS estimates, which we use in our final analysis.

In Figure 1, we report the results from a counterfactual experiment using the RIS estimates and the parametric bootstrap methods discussed in Section IV. We focus on Washington, Oregon, and Idaho, three states that adopted term limits in the early 1990's. Specifically, we are interested in what happens to the probability that Washington adopts term limits when we when we manipulate the underlying propensity of its neighbors—Oregon and Idaho—to adopt term limits. We start by taking 1,000 draws from a multivariate normal distribution with a mean equal to the vector of parameter estimates and a variance-covariance matrix equal to the estimated information matrix. For each draw, we calculate the probability that each state will adopt term limits when the y_i^* 's for Oregon and Idaho are set 1-unit above their estimated values, using the parameter draws and the observed values for the independent variables, and the y_i^* 's for the other forty-six states are held fixed at their estimated values (i.e., not manipulated). We then calculate these same probabilities when the y_i^* 's for Oregon and Idaho are set 1-unit below their estimated values and take the difference between the first and second vector of probabilities for each set of parameter draws. Figure 1 provides the empirical distribution of the 1000 changes in Washington's probability of adopting term limits given our counterfactual changes to Oregon and Idaho.

[Figure 1 Here]

Figure 1 clearly demonstrates that, for the overwhelming majority of trials, the predicted effect of decreasing the propensities that Oregon and Idaho will adopt term limits is to decrease the probability that Washington will adopt term limits. The median effect from our trials is -.09, and a

90% confidence interval runs from $-.289$ to $-.009$. Of course, it is possible to conduct similar experiments for any cluster of neighboring states (i.e., change the subjects in our experiment) and to counterfactually manipulate the presence or absence of direct democracy (i.e., change the treatment), but we save these exercises for later. In the end, our results, using the RIS estimator, suggest that states' decisions to adopt term limits are influenced by the experiences of their neighbors.

VII. Conclusion

Spatial interdependence is substantively and theoretically ubiquitous and important across social-science binary-outcomes. Standard ML-estimation of binary-outcome models in the presence of spatial interdependence are badly misspecified if that interdependence is ignored, but they are also misspecified (we suspect less badly, but we have not explored that systematically as yet), if that interdependence is reflected by inclusion of an endogenous spatial lag as an explainer. Spatial-lag probit models are difficult and highly computationally demanding, but not impossible, to estimate with appropriate estimators. We expect future work to demonstrate more fully and clearly the conditions under which expending such effort in estimation merits the gains. It is also possible, and by simulation feasible, to calculate and present properly the estimated spatial *effects* (as opposed to merely probit *coefficients*) on binary outcomes, along with their associated estimates of certainty. Doing so will be of great substantive advantage under any conditions, regardless of whether the sophisticated spatial-probit estimators offer much gain (in terms of bias, efficiency, or standard-error accuracy) from standard-probit estimators with spatial lags pretended to be exogenous.

REFERENCES

- Autant-Bernard, C. 2006. "Where Do Firms Choose to Locate Their R&D? A Spatial Conditional Logit Analysis on French Data," *European Planning Studies* 14(9):1187-1208.
- Banerjee, S., Carlin, B.P., Gelfand, A.E. 2004. *Hierarchical Modeling and Analysis for Spatial Data*. Boca Raton: Chapman & Hall.

- Beron, K.J., Murdoch, J.C., Vijverberg, W.P.M. 2003. "Why Cooperate? Public Goods, Economic Power, and the Montreal Protocol," *Review of Economics and Statistics* 85(2):286-97.
- Beron, K.J., Vijverberg, W.P.M. 2004. "Probit in a Spatial Context: A Monte Carlo Analysis," in L. Anselin, R.J.G.M. Florax, & S.J. Rey, eds., *Advances in Spatial Econometrics: Methodology, Tools and Applications*. Berlin: Springer-Verlag.
- Bhati, A.S. 2005. "Modeling Count Outcomes with Spatial Structures: An Information-Theoretic Approach," unpublished: Justice Policy Center, The Urban Institute.
<http://www.american.edu/cas/econ/faculty/golan/Papers/Papers05/BhatiPaper.pdf>, or
<http://www.uni-kiel.de/ifw/konfer/spatial/bhati.pdf>.
- Bolduc, D., Fortin, B., Gordon, S. 1997. "Multinomial Probit Estimation of Spatially Interdependent Choices: An Empirical Comparison of Two New Techniques," *International Regional Science Review* 20:77-101.
- Cain, B.E., Levin, M.A. 1999. "Term Limits." *Annual Review of Political Science* 2:163-88.
- Carey, J.M, Niemi, R.G., Powell, L.W. 1998. "The Effects of Term Limits on State Legislatures." *Legislative Studies Quarterly* 23(2): 271-300.
- Case, A. 1992. "Neighborhood Influence and Technological Change," *Regional Science and Urban Economics* 22:491-508.
- Chamberlain, G. 1993. "Feedback in Panel Data Models," *Harvard Institute of Economic Research Working Paper* #1656.
- Cho, W.T., Rudolph, T. 2007. "Emanating Political Participation: Untangling the Spatial Structure behind Participation," *British Journal of Political Science* 37(1):
- Coughlin, C.C., Garrett, T.A., Hernández-Murillo, R. 2003. *Spatial Probit and the Geographic Patterns of State Lotteries*. St. Louis Federal Reserve Bank Working Paper 2003-042B.
<http://research.stlouisfed.org/wp/2003/2003-042.pdf>.
- Darmofal, D. 2007. "Bayesian Spatial Survival Models for Political Event Processes," Unpublished:
<http://people.cas.sc.edu/darmofal/DarmofalBayesianSpatialSurvival.pdf>.
- Dubin, R.A. 1997. "A Note on the Estimation of Spatial Logit Models," *Geographical Systems* 4(2):181-93.
- Fleming, M.M. 2004. "Techniques for Estimating Spatially Dependent Discrete-Choice Models," in L. Anselin, R.J.G.M. Florax, & S.J. Rey, eds., *Advances in Spatial Econometrics: Methodology, Tools and Applications*. Berlin: Springer-Verlag.
- Franzese, R., Hays, J. 2003. "Modeling Spatial Relationships in International and Comparative Political Economy: An Application to Globalization and Capital Taxation in Developed Democracies," presented at the Annual Meetings of the Midwest Political Science Association.
- Franzese, R., Hays, J. 2004a. "Empirical Modeling Strategies for Spatial Interdependence: Omitted-Variable vs. Simultaneity Biases," presented at the 21st Summer Meeting of the Society for Political Methodology.
- Franzese, R., Hays, J. 2004b. "Modeling International Diffusion: Inferential Benefits and Methodological Challenges, with an Application to International Tax Competition," *Wissenschaftszentrum-Berlin SP II 2004 – 12*.

- Franzese, R., Hays, J. 2005a. "Modeling Spatial Interdependence in Comparative and International Political Economy with an Application to Capital Taxation," presented at the annual Meeting of the Midwest Political Science Association.
- Franzese, R., Hays, J. 2005b. "Spatial Econometric Modeling, with Application to Employment Spillovers and Active-Labor-Market Policies in the European Union," presented at Groningen University, workshop on 'Partisan Politics, Political Autonomy, and Policy Harmonization across Europe'.
- Franzese, R., Hays, J. 2006a. "Calculating and Presenting Spatio-Temporal Patterns of Diffusion and Interdependence," presented at the General Conference of the European Consortium of Political Research, Pisa.
- Franzese, R., Hays, J. 2006b. "Spatiotemporal Models for Political-Science Panel and Time-Series-Cross-Section Data," presented at the 23rd Annual Summer Meetings of the Society for Political Methodology.
- Franzese, R., Hays, J. 2006c. "Strategic Interaction among EU Governments in Active-Labor-Market Policymaking: Subsidiarity and Policy Coordination under the European Employment Strategy," *European Union Politics* 7(2):167-89.
- Franzese, R., Hays, J. 2007a. "Empirical Models of International Capital-Tax Competition," in *International Taxation Handbook*, G. Gregoriou, C. Read, eds., Elsevier Press, pp. 43-72.
- Franzese, R., Hays, J. 2007b. "Interdependence in Comparative & International Political Economy, with Applications to Economic Integration and Strategic Fiscal-Policy Interdependence," presented at Paris 13 (Université Paris), Axe 5: PSE.
- Franzese, R., Hays, J. 2007c. "Spatial-Econometric Models of Cross-Sectional Interdependence in Political-Science Panel and Time-Series-Cross-Section Data," *Political Analysis* 15(2):140-64.
- Franzese, R., Hays, J. 2007d. *Spatial Econometric Models of Interdependence*. Book prospectus.
- Franzese, R., Hays, J. 2008a. "Empirical Models of Spatial Interdependence," J. Box-Steffensmeier, H. Brady, D. Collier, eds., *Oxford Handbook of Political Methodology*, Oxford UP.
- Franzese, R., Hays, J. 2008b. "Interdependence in Comparative Politics: Substance, Theory, Empirics, Substance," *Comparative Political Studies* 41(4/5):742-80.
- Franzese, R., Hays, J. 2009a. "Contagion, Common Exposure, and Selection: Empirical Modeling of the Theories and Substance of Interdependence in Political Science," *Concepts & Methods: Newsletter of the International Political Science Association* .
- Franzese, R., Hays, J. 2009b. "Empirical Modeling of Spatial Interdependence in Time-Series Cross-Sections," in S. Pickel, G. Pickel, H-J. Lauth, D. Jahn, eds., *Neuere Entwicklungen und Anwendungen auf dem Gebiet der Methoden der vergleichenden Politikwissenschaft* (Methods of Comparative Political Science and Social Science: New Developments and Applications), Wiesbaden: Westdeutscher Verlag, pp. 233-261.
- Franzese, R., Hays, J., Kachi, A. 2009. "The m-STAR Model as an Approach to Modeled, Dynamic, Endogenous Interdependence in Comparative & International Political Economy." Presented at the Networks in Political Science (NIPS) Conference, Harvard Kennedy School of Government (2008); update presented at the Southern Political Science Association Conference (2009).

- Garrett, T.A., Wagner, G.A., Wheelock, D.C. 2005. "A Spatial Analysis of State Banking Regulation," *Papers in Regional Science* 84(4):575-95.
- Graham, E., Shipan, C.R., Volden, C. 2008. "The Diffusion of Policy Diffusion Research," unpublished, The Ohio State University and The University of Michigan.
- Hays, J.C., Kachi, A. 2008. "Estimating Interdependent Duration Models with an Application to Government Formation and Survival," plenary session paper given at the 25th Annual Summer Conference of the Society for Political Methodology, University of Michigan. Available at <http://polmeth.wustl.edu/conferences/methods2008/program.html>.
- Heckman, J.J. 1978. "Dummy Endogenous Variables in a Simultaneous Equation System," *Econometrica* 46:931-59.
- Holloway, G., Shankar, B., Rahmanb, S. 2002. "Bayesian Spatial Probit Estimation: A Primer and an Application to HYV Rice Adoption," *Agricultural Economics* 27(3):383-402.
- Honore, B.E., Kyriazidou, E. 2000. "Panel Data Discrete Choice Models with Lagged Dependent Variables," *Econometrica* 68(4):839-74.
- Hunter, D.R., Handcock, M.S., Butts, C.T., Goodreau, S.M., Morris, M. 2008. "ergm: A Package to Fit, Simulate and Diagnose Exponential-Family Models for Networks." *Journal of Statistical Software* (URL <http://www.jstatsoft.org/v24/i03/>).
- Kayser, M.A. 2007. "Partisan Waves: International Sources of Electoral Choice," unpublished. University of Rochester. http://mail.rochester.edu/~mksr/papers/PWaves_ECM_070108.pdf.
- Kelejian, H.H., Prucha, I.R. 2001. "On the Asymptotic Distribution of the Moran I Test Statistic with Applications," *Journal of Econometrics* 104:219-57.
- Lacombe, D.J., Shaughnessy, T.M. 2005. "An Examination of a Congressional Vote Using Bayesian Spatial Probit Techniques." Paper presented at the 2005 Meetings of the Public Choice Society.
- LeSage, J.P. 1999. *Spatial Econometrics*. <http://www.rri.wvu.edu/WebBook/LeSage/spatial/wbook.pdf>
- LeSage, J.P. 2000. "Bayesian Estimation of Limited Dependent Variable Spatial Autoregressive Models," *Geographical Analysis* 32(1):19-35.
- Lin, G. 2003. "A Spatial Logit Association Model for Cluster Detection," *Geographical Analysis* 35(4):329-40.
- McMillen, D.P. 1992. "Probit with Spatial Autocorrelation," *Journal of Regional Science* 32:335-48.
- McMillen, D.P. 1995. "Selection Bias in Spatial Econometric Models," *Journal of Regional Science* 35(3):417-36.
- Morehouse, S.M., Jewell, M.E. 2004. "States As Laboratories: A Reprise." *Annual Review of Political Science* 7: 177-203.
- Murdoch J.C., Sandler, T., Vijverberg, W.P.M. 2003. "The Participation Decision versus the Level of Participation in an Environmental Treaty: A Spatial Probit Analysis," *Journal of Public Economics* 87(2):337-62.
- Novo, A. 2003. *Contagious Currency Crises: A Spatial Probit Approach*. Banco de Portugal Working Paper: <http://www.bportugal.pt/publish/wp/2003-5.pdf>.

- Phaneuf, D.J., Palmquist, R.B. 2003. "Estimating Spatially and Temporally Explicit Land Conversion Models Using Discrete Duration," Unpublished:
http://www.aere.org/meetings/0306workshop_Phaneuf.pdf.
- Pinkse, J., Slade, M.E. 1998. "Contracting in Space: An Application of Spatial Statistics to Discrete-Choice Models," *Journal of Econometrics* 85: 125-54.
- Pinkse, J. 1999. *Asymptotic Properties of Moran and Related Tests and Testing for Spatial Correlation in Probit Models*. University of British Columbia, Department of Economics Working Paper:
- Pluemper, T., Neumayer, E. 2008. "Model Specification in the Analysis of Spatial Dependence," unpublished.
- Rathbun, S.L., Fei, S. 2006. "A Spatial Zero-Inflated Poisson Regression Model for Oak Regeneration," *Environmental and Ecological Statistics* 13(4):409-26.
- Robins, G, Pattison, P., Kalish, Y., Lusher, D. 2007. "An introduction to exponential random graph (p*) models for social networks," *Social Networks* 29(2):173-91.
- Schofield, N., Miller, G., Martin, A. 2003. "Critical Elections and Political Realignments in the USA: 1860-2000," *Political Studies* 51(2):217-40.
- Signorino, C. 1999. "Strategic Interaction and the Statistical Analysis of International Conflict," *American Political Science Review* 93(2):279-98.
- Signorino, C. 2002. "Strategy and Selection in International Relations," *International Interactions* 28:93-115.
- Signorino, C., Yilmaz, K. 2003. "Strategic Misspecification in Regression Models," *American Journal of Political Science* 47(3):551-66.
- Signorino, C. 2003. "Structure and Uncertainty in Discrete Choice Models," *Political Analysis* 11(4): 316-44.
- Signorino, C., Tarar, A. 2006. "A Unified Theory and Test of Extended Immediate Deterrence," *American Journal of Political Science* 50(3):586-605.
- Simmons, B., Elkins, Z. 2004. "The Globalization of Liberalization: Policy Diffusion in the International Political Economy," *American Political Science Review* 98(1):171-89.
- Smith, T.E., LeSage, J.P. 2004. "A Bayesian Probit Model with Spatial Dependencies," in J.P. LeSage & R.K. Pace, eds., *Spatial and Spatio-Temporal Econometrics*. Amsterdam: Elsevier.
- Snijders, T.A.B. 2005. "Models for Longitudinal Network Data." In P. Carrington, J. Scott and S. Wasserman, eds., *Models and Methods in Social Network Analysis*. New York: Cambridge.
- Snijders, T.A.B., Steglich, C., Schweinberger, M. 2007. "Modeling the Co-Evolution of Networks and Behavior," in K. van Montfort, H. Oud, A. Satorra, eds., *Longitudinal Models in the Behavioral and Related Sciences* (Mahwah NJ: Lawrence Erlbaum), pp.41-71.
- Volden, Craig. 2006. "States as Policy Laboratories: Emulating Success in the Children's Health Insurance Program." *American Journal of Political Science* 50(2): 294-312.

Table 2: Term Limits and Direct Democracy

State	Term-Limits (Year)	Repealed (Year)	Ballot Initiatives	Popular Referenda
Arizona	Yes (1992)	No	Yes	Yes
Arkansas	Yes (1992)	No	Yes	Yes
California	Yes (1990)	No	Yes	Yes
Colorado	Yes (1990)	No	Yes	Yes
Florida	Yes (1992)	No	Yes	No
Idaho	Yes (1994)	Yes (2002)*	Yes	Yes
Illinois	No	—	Yes	Yes
Kentucky	No	—	No	Yes
Louisiana	Yes (1995)	No	No	No
Maine	Yes (1993)	No	Yes	Yes
Maryland	No	—	No	Yes
Massachusetts	Yes (1994)	Yes (1997)	Yes	Yes
Michigan	Yes (1992)	No	Yes	Yes
Mississippi	No	—	Yes	No
Missouri	Yes (1992)	No	Yes	Yes
Montana	Yes (1992)	No	Yes	Yes
Nebraska	Yes (2000)	No	Yes	Yes
Nevada	Yes (1996)	No	Yes	Yes
New Mexico	No	—	No	Yes
North Dakota	No	—	Yes	Yes
Ohio	Yes (1992)	No	Yes	Yes
Oklahoma	Yes (1990)	No	Yes	Yes
Oregon	Yes (1992)	Yes (2002)	Yes	Yes
South Dakota	Yes (1992)	No	Yes	Yes
Utah	Yes (1994)	Yes (2003)*	Yes	Yes
Washington	Yes (1992)	Yes (1998)	Yes	Yes
Wyoming	Yes (1992)	Yes (2004)	Yes	Yes

Notes: In Idaho and Utah, term limits were repealed by their state legislatures (*). Term-limits were overturned by state supreme courts in MA, OR, WA, and WY. Our sample only includes the contiguous 48 states. Alaska allows both ballot initiatives and popular referenda, but has never adopted term limits. Hawaii allows neither ballot initiatives nor popular referenda and has never adopted term limits.

Table 3: Adoption of Term Limits for State Legislators, Estimation Results

	Probit-ML	Probit-MCMC	Spatial-Lag Probit	Spatial-Error Probit	Probit-RIS
<i>Constant</i>	-.539 (2.579)	-.909 (1.814)	-.598 (3.080)	-.411 (2.533)	.313 (2.223)
<i>I&R</i>	2.320*** (.581)	1.806*** (.481)	3.257*** (.917)	2.650*** (.627)	2.336*** (.620)
<i>Clinton</i>	.273 (.542)	.131 (.476)	.147 (.769)	.056 (.606)	.304 (.518)
<i>Tax Effort</i>	-.178 (.273)	-.068 (.183)	-.176 (.321)	-.146 (.262)	-.214 (.247)
<i>Spatial lag or error-lag</i>	.926 (.801)	.634 (.687)	.144 (.207)	.018 (.279)	.416** (.194)
Pseudo-R²	.480	.458	.833	.803	—
Log-Likelihood	-17.093	—	—	—	-16.261
Observations	48	48	48	48	48

Notes: The first two columns' models are estimated assuming the spatial lags exogenous. The first column estimates are from the standard ML estimator. Its parentheses contain estimated standard errors; its hypothesis tests assume asymptotic normality of calculated *t*-statistics. The models in columns two through four apply MCMC methods with diffuse uninformative priors. The reported coefficient estimates are the posterior-density means based on 10,000 samples after 1000-sample burn-ins. The parentheses contain sample standard-deviations of these posteriors. The *p*-values are calculated directly from the posterior density without calculating *t*-statistics or assuming normality. ****p*-value <.01, ***p*-value<.05, **p*-value <.10.

Figure 1: *Parametric Bootstrap:* Counterfactual Effect on Probability Washington State Adopts Term Limits from 2-Unit Decreases (+1 to -1) in Oregon & Idaho's Propensities to Adopt Limits

