Knox meets Cox:
Adapting Epidemiological Space-Time Statistics to Demographic Studies*

Carl P. Schmertmann
Center for Demography and Population Health
Florida State University
Email: schmertmann@fsu.edu
Phone: 1-850-644-7100 Fax: 1-850-644-8818

Renato M. Assunção
Departamento de Estatística
Universidade Federal de Minas Gerais

Joseph E. Potter
Department of Sociology
University of Texas at Austin

* This research was supported by a grant from the National Institutes of Health, Institute of Child Health and Human Development (Grant No. R01 HD41528). Suzana Cavenaghi was a key collaborator in the earlier stages of our research, and Márcio Thomé provided us with historical information on family planning programs in Rio Grande do Norte. Data and *R* programs for reproducing our results are available online at *http://schmert.net/KnoxCox* .

ABSTRACT

Many important questions and theories in demography focus on changes over time, and on how those changes differ over geographic and social space. Space-time analysis has always been important in studying fertility transition, for example, but with few exceptions demographers have not used formal statistical methods to describe and analyze time series of maps. One formal method, used widely in epidemiology, criminology and public health, is Knox's space-time interaction test. In this paper, we discuss the potential of the Knox test in demographic research, and note some possible pitfalls. We demonstrate how to use familiar proportional hazards models to adapt the Knox test for demographic applications. These adaptations allow for non-repeatable events, and for the incorporation of structural variables that change in space and time. We apply the modified test to data on the onset of fertility decline in Brazil over 1960-2000, and show how the modified method can produce maps showing where and when diffusion effects seem strongest, net of covariate effects.

INTRODUCTION

Space-time analysis has long been central to the study of population. Many of the important questions and theories in demography focus on changes over time, and on how those changes differ over geographic and social space. Perhaps the most prominent example is the European Fertility Project (Coale and Watkins 1986), which compiled sequences of maps for various indices of fertility and nuptiality across European provinces. These maps, especially one displaying estimated dates of onset for sustained decline in marital fertility, served as an important underpinning for the argument that cultural factors and the diffusion of information or attitudes toward family limitation were critical to the onset of fertility decline.

Aided by advances in statistical methods and software, demographers are paying increasing attention to spatial patterns in population data. Several important papers have applied or adapted techniques from spatial econometrics (especially the ideas of Anselin 1988 and Land and Deane 1992) to study area-level data. Examples include Baller et al. (2001), Baller and Richarson (2002), Land, Deane, and Blau (1991), and Tolnay (1995).

With some important exceptions (Tolnay et al. 1996), spatial analysis in demography has omitted a time dimension and used purely cross-sectional statistics. In contrast, related fields such as epidemiology, criminology, and public health have integrated time more directly into statistical analyses of spatial event patterns. Recent research in these fields has used formal statistical methods, mainly from epidemiology, to study *sequences* of cross-sectional maps.

The lack of crossover of epidemiological statistics into demographic research has several origins. Disciplinary differences in training, combined with a relative lack of

longitudinal spatial data in demography (at least until recently) explain part of the gap. Standard methods in space-time statistics can be also ill suited to demographic studies, which often involve longer periods than epidemiological or crime studies. Many epidemiological methods are appropriate for short time scales, over which it is reasonable to assume that background variables such as the population at risk or local socioeconomic conditions are constant. Such assumptions are often unrealistic in demographic studies, making direct application of standard epidemiological methods difficult or questionable.

In this paper we try to bridge the gap between demographic questions and epidemiological methods. We begin by introducing a well established and still widely used formal method for studying space-time interactions in epidemiology -- the Knox test. We discuss its utility to demographers, and show how to adapt the Knox test in a specific demographic problem, looking for evidence of social contagion in fertility transition. These adaptations use Cox's (1972) proportional hazards framework to model covariate effects that may otherwise yield false positives in a search for epidemic patterns, and to model survival processes with non-repeatable events.

CONTAGION AND MAPPING

Demographers and epidemiologists share an interest in understanding infection and contagion. In demography, 'infection' can be interpreted broadly to represent not only disease, but also the spread of an idea or behavior from one person to another – for example, adoption of a new contraceptive method. Many social processes mimic contagious diseases, and there is a substantial demographic literature suggesting that interpersonal contacts and infectious ideas are important components of population-level demographic change (Montgomery and Casterline 1996).

What would a contagious process look like on a map? In a statistical sense, if a process is infectious then rates are history-dependent: occurrence of an event increases the hazards of future events, particularly in nearby places and in the near future. An important implication of such rate dependence is that events generated by epidemics (biological or social) will tend to cluster in space and time.

Spatial or temporal event clusters alone do not necessarily demonstrate contagion, however. A non-contagious process could have a strong spatial pattern if population is unevenly distributed over space, or if local factors imply differential risks across locations. Similarly, temporal trends in population or other risk factors could generate an uneven distribution of events over time.

The hallmark of contagion is second-order, or space-time, clustering. A contagious process will generate a changing spatial distribution of incidence, with event clusters in different places at different times. *Space-time interaction* in event data is a key indicator of an infectious pattern, and it forms the basis of statistical tests for contagion.

Figure 1 displays estimated fertility transition times for 502 regions of Brazil over 1960-2000. Each panel corresponds to an intercensal period, with shaded areas representing regions undergoing transitions in that period. In Figure 1 and throughout this paper, we define a transition as either a 20% or greater decline in the total fertility rate (TFR), or a TFR below 4.5 at the end of the period. We discuss these data more carefully in later sections, but here they serve an example of strong space-time interaction. Note the clear temporal clustering (very few events before 1960 and after 1991), but also the clear space-time clustering (very different spatial distributions in the six periods). Strong space-time interaction is consistent with a contagious process, and indeed demographers

have used maps like Figure 1 as evidence for infectious spread of new fertility norms through social interactions (Coale 1973).

[FIGURE 1 about here: transition times]

Data such as those in Figure 1 do not make an open-and-shut case for the epidemic nature of demographic events, however. A skeptical demographer considering Figure 1 as evidence of 'infection' would have some serious reservations.

Most importantly, the time period is longer than those in epidemiological applications. When examining events over periods of weeks or months, it is reasonable to assume that factors affecting relative risks across areas (resident populations, socioeconomic conditions, age structures, etc.) are approximately constant. In that case a changing spatial pattern of events over time, like that in Figure 1, provides good evidence for contagion.

Our skeptical demographer would note, however, that constant relative risks are an extremely poor assumption for demographic processes like fertility transition. Demographic events often unfold over long periods, against a background of large changes in underlying conditions that differentially affect the risks of events in different locations. Space-time interaction in demographic events is evidence for contagion, in short, only if it is unexplained by space-time patterns in other relevant variables.

KNOX'S FORMAL TEST FOR CONTAGION

Event Maps, Event Clusters, and the Knox Statistic

We begin with notation. Consider a demographic process over a space **S** and time interval **T**. In most applications **S** would be either a discrete set of spatial units such as {502 regions in Brazil}, or a continuous set of points such as {all (Latitude, Longitude)

pairs in Brazil}. Similarly **T** could be either a set of discrete times such as census dates, or a continuous interval such as [1960,2000].

We observe *n* events, *i*=1…*n*, each with a known location $s_i \in \mathbf{S}$ and time $t_i \in \mathbf{T}$. One can create a drawing to indicate the *n* places and times, as in Figure 1, but for statistical purposes it is useful to define a *map* more abstractly, as a set:

$$M = \{(s_i, t_i)\}_{i=1\ldots n}$$

Index numbers for events *i*=1…*n* are arbitrary because *M* is an unordered set. In this paper we follow the convention of numbering events in chronological order, so that $t_1 \le t_2 \le \cdots \le t_n$.

Measuring event clustering requires definitions of spatial and temporal proximity. Following Knox (1964a, 1964b), construct dummy indicators for each pair of events (*i*,*j*):

$$(1) \qquad \sigma_{ij} = I\left[\, s_i, s_j \text{ are close in space} \,\right], \quad \tau_{ij} = I\left[\, t_i, t_j \text{ are close in time} \,\right]$$

where the indicator function I[ ] equals one if the condition in brackets is true and zero otherwise. In general, researchers using the Knox test must define 'closeness' in space and time in the context of the particular study.[1] In this paper we use a slight variant of the definitions used by Bocquet-Appel and Jakobi (1998) in a British fertility study: we define transitions in Figure 1 as close in space if they occurred in adjacent regions or in regions with centroids <100km apart, and as close in time if they occurred in the same intercensal period.

[1] Examples in applied studies include [≤1 km apart, ≤59 days] in Knox's (1964b) analysis of childhood leukemia, [<5 km, <1 year] by McNally et al. (2002) in a study of childhood brain tumors, [< 600 meters, <7 days] in Knox's (2002) study of serial murders, [≤5 meters, ≤3 days] in Morisson et al.'s (1998) study of a dengue fever outbreak in a Puerto Rican town, and [≤100 km, "same intercensal period"] in Bocquet-Appel and Jakobi's (1998) analysis of British fertility transitions.

Given definitions for spatial and temporal proximity, for n events on a map there are $N=n(n-1)/2$ distinct event pairs, of which

$$(2) \qquad N_s = \sum_i \sum_{j>i} \sigma_{ij} \quad , \quad N_t = \sum_i \sum_{j>i} \tau_{ij} \quad , \quad X = \sum_i \sum_{j>i} \sigma_{ij}\ \tau_{ij}$$

are respectively close in space, close in time, or close in both. The value of X is the map's *Knox statistic*, which is central to testing for contagion. X is a simple index of space-time clustering, and in a contagious process it should be high relative to the level of spatial and temporal clustering (N_s and N_t, respectively).

As an instructive example, consider a map for n=10 new cases of a disease occurring over five possible locations **S**={A,B,C,D,E} at three possible time periods **T**={1,2,3}:

M = { (E,1), (A,2), (A,2), (A,2), (A,3), (B,3), (C,3), (C,3), (C,3), (D,3) }

For this map the first event was in location s_1=E at time t_1=1, the second was in s_2=A at t_2=2, and so on. There are n=10 events and 45 distinct event pairs. If we define events as spatially close when they occur in the same location, then N_s=9 for this map, because event pairs 2&3, 2&4, 2&5, 3&4, 3&5, 4&5, 7&8, 7&9, and 8&9 share locations. Analogously, if events are temporally close when they occur in the same period, then N_t=18. Under these definitions, X=6 of the 45 event pairs (2&3, 2&4, 3&4, 7&8, 7&9, 8&9) are close in both space and time.

Permutation Maps

The main question for detecting contagion is whether the observed space-time clustering in a map is high, given the levels of spatial and temporal clustering. For example, is X=6 high for a map with n=10, N_s=9, and N_t=18? One way to answer this

question is by comparing the Knox statistic X for the realized map (M) to the Knox statistic for another map (M^*) with the same event locations $\{s_1 \ldots s_n\}$ and times $\{t_1 \ldots t_n\}$, but different *pairings* of locations and times. If the observed map M has a higher Knox statistic X than most random pairings, then there is evidence of space-time event clustering, and contagion is more likely.

Generating an alternative map with identical $\{s_1 \ldots s_n\}$ and $\{t_1 \ldots t_n\}$ is straightforward. Each possible pairing of locations and times corresponds to a different permutation π of the indices $1 \ldots n$. Using the standard notation (i) to represent the i^{th} element of the permutation, one can construct $M^*(\pi)$, the alternative map for a permutation $\pi=[(1) \ldots (n)]$, by pairing reordered locations $\{s_{(1)} \ldots s_{(n)}\}$ with the original times $\{t_1 \ldots t_n\}$. Intuitively, a permutation map is simply a reshuffling of the n event locations, holding the event times constant.[2] For instance, one possible permutation of our example map is $\pi=[4,8,1,5,10,2,9,7,6,3]$, in which case the reordered locations are

$$\{ s_{(1)}, s_{(2)}, s_{(3)} \ldots s_{(10)} \} = \{ s_4, s_8, s_1 \ldots s_3 \} = \{ \text{A, C, E, A, D, A, C, C, B, A} \}$$

and the reordered map is

$$M^*(\pi) = \{ \text{(A,1), (C,2), (E,2), (A,2), (D,3), (A,3), (C,3), (C,3), (B,3), (A,3)} \}$$

Like the original M, the permutation map $M^*(\pi)$ has $N_s=9$ and $N_t=18$. These equalities happen by construction, because both maps have the same marginal distributions of event locations and event times. Unlike the $X=6$ original map, however, space-time clustering in $M^*(\pi)$ is only $X^*(\pi)=2$. This particular π is only one of $10! > 3.6$

[2] Pairing reordered locations with original times is arbitrary. The reverse process, with original locations and reordered times, produces an identical set of permutation maps. Keeping times in chronological order simplifies notation when discussing survival processes, because $(1) \ldots (n)$ represents the order in which failures happen – first location $s_{(1)}$, then $s_{(2)}$, and so on.

million possible permutations, but if this result is typical it suggests that map *M* has a high level of space-time clustering.

The Knox Test

The Knox test formalizes the comparison with permutation maps, by considering the distribution of the space-time clustering statistic $X^*(\pi)$ across all possible pairings of observed event times and places. The strength of evidence for space-time clustering can be represented by a *p*-value for the observed Knox statistic *X*:

$$(3) \qquad p_0(X) = \Pr\left[X^* \textit{ for a random permutation} \geq X \right] = \sum_{\pi} L_0(\pi)\, I\left[X^*(\pi) \geq X\right]$$

where the sum is over all $n!$ possible permutations π, and $L_0(\pi)$ is the probability of particular permutation under the null hypothesis of no contagion. A small value such as $p < .01$ means that the Knox statistic *X* is unusually high and the observed map gives strong evidence for space-time interaction in event rates.

The standard Knox test derives the *p*-value in Equation 3 by imposing a strong mathematical assumption. Specifically, one must assume that an absence of contagious effects implies *constant relative risks* (CRR). CRR requires that all local risks change by the same multiplicative factors at each moment. For example, if event risk is twice as high in location *A* as in *B* at the start of the study period, it must remain twice as high at all times during the study. As a corollary to CRR, if locations *A* and *B* both have positive event risks at time $t=0$, then there can be no later time at which *A* has zero risk and *B* has positive risk.

The CRR assumption guarantees that all map permutations π are equally likely in the absence of contagion (see Appendix), so that $L_0(\pi)=1/n!$ for every permutation and Equation 3 becomes

$$(4) \qquad p_0(X) = \frac{1}{n!} \sum\nolimits_{\pi} I\left[X^*(\pi) \geq X\right]$$

Analysts have used Equation 4 to perform Knox tests in two ways. First, the equation allows elegant analytical approximations to the null distribution of X under CRR (Barton and David 1966; Kulldorff and Hjalmars 1999). More importantly for our purposes, it also makes Monte Carlo simulation easy. Although in practice *n!* is prohibitively large (e.g., 10! > 3.6 million, 100! > 10^{158}), one can approximate the right-hand side of Equation 4 by drawing a large number K of permutations independently and uniformly (i.e., all permutations are equally likely) and approximating with

$$(5) \qquad \hat{p}_0(X) = \frac{1}{K} \sum_{k=1}^{K} I\left[X^*(\pi_k) \geq X\right]$$

This standard Monte Carlo estimate, typically performed over several thousand random permutations, is simply the proportion of sampled permutation maps with Knox statistics at least as large as X. A small $\hat{p}_0(X)$ means that the observed events have unusually high space-time clustering, thus providing evidence for an infectious process.

Although it is strong, the CRR assumption is reasonable for repeatable-event processes like infections observed over short periods. However, the CRR assumption means that the standard Knox test is inconsistent with either (1) a demographic process in which changing covariates alter the relative risks in locations over time, or (2) a survival process in which events (such as fertility transitions) remove discrete locations from the risk set. We return to these concerns shortly.

Applications of the Standard Knox test

Researchers have used the standard Knox test for many purposes. Bhopal *et al.* (1992) applied the test to find strong space-time interaction in cases of Legionnaire´s disease in Glasgow and Edinburgh. Samuelsson *et al.* (1994) detected space-time clustering in Swedish data, suggesting that infectious agents are partially responsible for insulin-dependent diabetes. Birch *et al.* (2000) reached the same conclusion for childhood leukemia. McKenzie *et al.* (2005) found evidence for epidemic patterns in suicides among mental health patients. Machado-Coelho et al. (1999) found significant space-time clustering in leishmaniasis incidence in Southeastern Brazil. The Knox test has also been used to show lack of support for associations between environmental hazards and health conditions, as in Morris *et al.*'s (1998) analysis of reported space-time clusters of Down's syndrome in England and Wales.

Several authors have suggested using Knox tests and space-time interaction patterns as early-warning systems. These efforts include Knox's (2002) analysis of space-time clusters in murders committed by a serial killer, and Tobin's (2007) study of the spread of an invasive gypsy moth in the US Midwest.

In demography, the most important applications of the Knox test have been in studying fertility transitions. Bocquet-Appel and Jakobi (1998) analyzed the geographical pattern of fertility change in 78 counties of Great Britain between 1861 and 1901. They generated a sequence of smoothed maps of the I_g fertility index for the five census years between 1861 and 1901, and from those maps estimated the time at which each county began its fertility transition. They used standard Knox tests to investigate spatial diffusion of fertility control, and found strong evidence of space-time interaction in onset times

(p=.002 or .001, depending on the definition of spatial proximity). Thus, nearby counties tended to have similar onset times, as in an infectious process. Bocquet-Appel and Jakobi (1998) concluded that fertility decline in Britain had characteristics of a social epidemic.

Applying a similar statistical test to India, with data for several hundred districts from 1961 to 1991, Bocquet-Appel and colleagues concluded that the fertility transition there did not exhibit strong epidemicity (Balabdoaoui *et al*. 2001; Bocquet-Appel *et al*. 2002). Because space-time interaction is weak in the Indian maps, these researchers attribute fertility decline in India mainly to non-contagious processes, such as nationwide family planning policies.

ADAPTING THE KNOX TEST FOR DEMOGRAPHIC APPLICATIONS

While some analysts have used the standard Knox test in demographic studies, its assumptions are ill suited to processes that evolve over long periods with changing risks. In such processes the standard test is likely to bias the approximation in Equation 5 downwards, making it too easy to conclude that a process is contagious.

Simply put, the standard Knox test attributes to contagion *any* space-time interactions in the data, even those generated by changing risk sets or by changing local covariates. In either of these situations the CRR assumption is violated and the standard Knox test becomes more likely to yield a false positive for contagion.

In the remainder of this paper we study a survival process for a set of discrete locations, each of which has at most one event. Such a process violates the CRR assumption because the risk set changes as the process unfolds, thus altering the relative risks of events in different locations (i.e., the risk of further events becomes zero once a location 'fails', while remaining positive in other locations).

Even with these new complications, the essential logic of the Knox test remains valid: if observed X is in the upper tail of values for maps with the given spatial and temporal frequencies of events, then the map provides evidence of a contagious process. The novelty in survival or covariate models is that the CRR simplification does not hold. The absence of space-time interaction no longer implies that all event orders π are equally likely in the absence of contagion. In such models we can no longer sample uniformly from the $n!$ permutations to generate simulated maps and X^* values under the null.

Fortunately, the statistical theory of event order in survival models with covariates is well understood from the literature on proportional hazards models. If each of n units has a constant failure hazard of ρ_i and there is no contagion, then the probability that units fail in order $\pi = (1)\ldots(n)$ is

$$(6) \qquad L_0(\pi) = \prod_{i=1}^{n} \left(\rho_{(i)} \Big/ \sum_{j \in R_{(i)}} \rho_j \right)$$

where $R_{(i)}=\{(i)\ldots(n)\}$ represents the set of units still at risk for the ith failure. By extension, if hazards change over time and are affected by local covariates z, but satisfy the proportionality assumptions of Cox (1972):

$$(7) \qquad \rho(s,t) = \gamma(t) \exp\left[\beta' z(s,t)\right]$$

then without contagion the probability of failures in order $\pi = (1)\ldots(n)$ is

$$(8) \qquad L_0(\pi \mid \beta) = \prod_{i=1}^{n} \left(\exp\left\{\beta' z[s_{(i)}, t_i]\right\} \Big/ \sum_{j \in R_{(i)}} \exp\left\{\beta' z[s_j, t_i]\right\} \right)$$

which is identical to Cox's partial likelihood function (cf. Kalbfleisch and Prentice 1973, equation 2).

To extend the Knox approach to survival processes with covariates, we use Equation (8) in two ways. First, we find the covariate effects $\hat{\beta}$ that maximize the likelihood of the *observed* event order on the map. This is a standard proportional hazards regression. Second, we use the estimated $\hat{\beta}$ to draw *simulated* orders π from the non-uniform distribution implied by Equation (8).

This new procedure models the covariate effects in the Knox statistic X, and compares the observed statistic to a null distribution of X^* values that are similarly affected by covariates. The modified test allows us to build into the Knox simulation the predicted effects of covariates on (s,t) pairings. By drawing π values from an appropriate non-uniform distribution, we weight draws of $X^*(\pi)$ toward pairings of event locations and times that are more likely to occur, given the observed patterns in covariates. The result is a test that is better adapted than the standard Knox test to demographic data, with lower bias and lower probabilities of false positives for contagion. Note that the standard Knox test ($\beta=0$; no covariate effects) is a special case of the modified test.

DATA

As an illustration, we use Brazilian Demographic Census microdata from 1960, 1970, 1980, 1991, and 2000 to calculate fertility rates and socioeconomic indicators for each of 502 subnational regions on each census date. The regions have unchanging boundaries, and are identifiable from the geographic codes on all five censuses. Data come from long-form questionnaires that collected information on births to women of reproductive age. Potter, Schmertmann, and Cavenaghi (2002) provide a full description of the estimation procedures.

We defined pairs of regions as close in space if they were adjacent on the map, or if the distance between centroids was less than 100 km. Brazil's continental scale and uneven population distribution create some large and sparsely populated regions in the North and West. Our definition of spatial proximity compromises between adjacency/shared borders (which would omit from the 'close in space' pairs close-together but non-adjacent regions in the South) and using distance (which would omit adjacent regions in the North whose centroids are far apart. With 502 regions there are 125751 distinct pairs of regions, of which N_s=1797 are 'close in space' under this definition.

We classify a region as having had a fertility transition if either of two conditions is satisfied in a census year: TFR < 4.5, or TFR fell by more than 20% in the previous intercensal period. The typical definition of onset is a 10% fall in fertility (Coale and Treadway 1986; Guinnane, Okun, and Trussell 1994). We chose the stricter criterion mainly because we are using TFR estimates that are sometimes noisy for smaller, less-populated regions. These transition times were displayed earlier in Figure 1. Table 1 shows the distribution of transition times across periods, and Table 2 has example TFR data and transition times for nine regions.

[TABLE 1 about here: transition times]

[TABLE 2 about here: example transition data]

We define transition pairs as 'close in time' if they happened in the same intercensal period. This yields a temporal clustering measure of N_t=38184. The observed Knox statistic for the map is X=1119, as shown in Table 3.

[TABLE 3: Knox contingency table]

Applying the standard Knox permutation test to the transition map, we find extremely significant space-time clustering: of 10000 uniformly-drawn Monte Carlo reorderings, none produced Knox statistics $X^*(\pi)$ as large as 1119. Simulated X^* values ranged from 475 to 624, with a median of 545. The estimated mean and standard deviation of the X^* distribution were 545.4 and 19.4, respectively. The leftmost curve in Figure 2 illustrates the results from the standard permutation test, in the form of estimated p-values from Equation (5). The horizontal line at p=.05 intersects the curve at X=578, indicating that observed space-time clustering must exceed 578 to reject the null of no contagion at the 5% level. The observed Knox statistic X=1119 on the far right easily meets this criterion, and in fact far exceeds all 10000 simulated X^* values.

Results: Modifying the Knox Test for a Survival Process with Covariates

The standard test is biased in favor of finding contagion. It detects not only space-time interaction caused by an epidemic process, but also space-time interaction caused by changing risks and risk sets. To address this problem we use the proportional hazards correction. We use end-of-period covariates as predictors for transition hazards, allow covariates to vary within regions across periods, and treat the two post-2000 transitions as right-censored observations. We include five time-varying covariates – percent of households in the region with electricity, average education of adult women (in years), percent of women in the labor force, percent of the region's labor force (both sexes) in the primary sector, and percent Catholic – in a proportional hazards model. Table 4 shows the regression results.

[Table 4 about here: PH regression coefficients]

Our interest lies in the overall explanatory power of the model rather than the individual coefficients. In this regard the regression does well. Predicting that the 32 regions with the highest predicted risks would transition before 1960, that the 65 regions with the next-highest risk would transition during 1960-1970, and so forth, yields matching predicted and observed times for 325 of 502 regions (65%). In comparison, one would expect only 30.5% matching if simulated times were drawn in the same way without regard to predicted risk.

Because socioeconomic covariates are clearly related to transition times, the space-time event clusters that contribute to the Knox statistic (e.g., the concentration of early fertility transitions in the South and Southeast) may result from spatial clusters of covariates (e.g., greater electrification in the South and Southeast in early periods) rather than from a contagion. To investigate the effects of contagion, net of covariate effects, we use the modified Knox test. As in the standard test, the first 32 simulated transitions are assumed to occur before 1960, the next 65 between 1960 and 1970, etc. However, in the modified test the sampling of π is non-uniform, with covariate levels affecting each region's transition probability in each time period. This approach draws samples of discrete times that match the assumptions of the proportional hazards model, and also match the marginal distribution of t in the observed map. Matching both the model and the observed map is possible because the baseline hazard $\gamma(t)$ is unspecified in the proportional hazards model in Equation (7). The model therefore assigns a likelihood only to the order of events, with event times being arbitrary as long as they satisfy the ordering.

The modified Knox permutation test appears as the rightmost curve in Figure 2, together with the standard test reported earlier. Including covariates and shifting risk sets moves the null distribution higher by about 170 points. Mean X^* over 10000 random maps goes from 545 (standard test) to 706 (modified test). In statistical terms this is a massive shift: with covariate effects, the mean Knox value for a random map shifts upward by 8.2 of the original distribution's standard deviations (σ). Similarly, the Knox value required for a significance level of p=.05 rises from 578 to 753 (+9.0σ), and for p=.01 rises from 593 to 774 (+9.3σ).

The changes in Figure 2 illustrate two important points. First, standards of evidence for space-time interaction in the modified Knox test are much more stringent. Covariate effects could cause some of the space-time interaction captured by the standard Knox statistic, and the rightward shift of the curve shows that attributing all of the clustering to contagion would be a mistake. Second, the observed X=1119 is still highly significant in the modified test. In this data set, strong evidence remains against a null hypothesis of space-time independence in events, even after controlling for the included covariates and changing risk sets.

We add one important caution. Our example shows that including the effects of space-time patterns in relevant covariates can dramatically raise the critical value in the Knox test, thus making a false conclusion of contagion much less likely. However, false positives are still possible if the list of covariates is incomplete. Both the standard and modified Knox tests answer the question "Is there evidence of space-time interaction in events *beyond that explained by any included covariates?*". In the standard Knox test there are no covariates at all. In this paper we show how to account for covariate effects

that might be misinterpreted as contagion in the standard test. However, the same logical problem remains (albeit to a much lesser degree) in a modified Knox test: omitted covariate effects can still be misinterpreted as contagion. As in any regression model, failure to control for relevant covariates is a potential danger, and researchers should interpret test results within the context of the specific model.

Sensitivity Tests

The Knox statistic and the Knox tests depend on arbitrary definitions of spatial and temporal closeness. Our example also depends on an arbitrary definition of fertility transition and on a grainy set of only six possible transition times. In order to investigate the robustness of the example presented in Figure 2, we experimented with alternative definitions for fertility transition, and for the spatial and temporal proximity of transition pairs.

All definitions yield essentially the same conclusions regarding contagion, and regarding the importance of covariate effects. Table 5 reports results from 36 experiments. We used three possible transition definitions: [TFR < 4.5 at end of period or TFR decreased by at least 20% over the period], [TFR < 4.0 at end of period], or [TFR decreased by at least 10% over the period], in combination with six possible definitions of spatial proximity (third column) and two possible definitions of temporal proximity (fourth column).

[TABLE 5 here – Sensitivity Tests]

In each experiment we calculated the Knox statistic for the observed map under the given definitions (X). We then drew 1000 random values of X^* from the standard null

distribution, assuming uniform map permutation probabilities. Table 5 reports the standard deviation (σ), mean (μ), and 95th percentile for this standard distribution. Finally, for each set of definitions we drew another 1000 X^* values from the modified null distribution with covariate effects. The rightmost column of Table 5 reports the upward shift in the 95th percentile of X^* values, measured in standard deviations σ. Large shifts in this critical value represent much more stringent standards of sample evidence for rejecting the null hypothesis of no contagion.

The baseline experiment, using the definitions we eventually adopted, appears in the first row of Table 5. For this case we repeat the results from the 10000 Monte Carlo permutations already reported: the 95th percentile for the null distribution of X^* shifts from 578 to 753, an increase of 9.0 standard deviations. Table 5 shows that such results are typical. Although the level of X varies with the choice of proximity definitions, the modified Knox test always requires much stronger statistical evidence than the standard Knox test. Indeed, the smallest upward shift in the 5% critical value in any of our experiments (Cases 26 and 32) was 4.0 standard deviations, and the average shift in the 5% critical value was more than 10 standard deviations. In short, including effects of space-time patterns in covariates on the Knox statistic X appears to alter the Knox test radically, regardless of the definition of events, or of the specific definitions of spatial and temporal proximity between events.

Results: Early and Late Fertility Transitions

Our modified Knox test also provides a new way to analyze the map's space-time interactions. Specifically, we can identify which regions had unexpectedly early or late fertility transitions, considering their socioeconomic conditions. We do this by comparing

each region's observed transition time t_i to the distribution of its simulated transition times $\{t_i^*(\pi_k)\}$ k=1…10000 across Monte Carlo samples. Such comparisons, when mapped, provide information about where and when the spatial interactions detected by the modified Knox test occurred.

In order to study early and late transitions, we define indicators for each region:

$$E_i = I\left[\, t_i < t_i^*(\pi_k) \text{ in over 60\% of samples} \,\right]$$
$$L_i = I\left[\, t_i > t_i^*(\pi_k) \text{ in over 60\% of samples} \,\right]$$

E_i=1 if an area's transition occurred earlier than one would have predicted from its covariate levels; L_i=1 if the transition was unexpectedly late.[3] Using this definition, 86 of 502 transitions were early, 37 were late, and 379 were neither early nor late. Table 6 shows the complete distribution of early and late transitions across regions and intercensal periods.

[Table 6 about here]

As examples, Table 7 below shows the distributions of simulated times t^* for the nine regions shown earlier in Table 2. Each row of Table 7 contains the region's distribution of simulated (covariate-adjusted) transition times across periods. Observed transition times for each region from Table 2 are indicated by shaded cells. The last row of Table 7 shows the marginal distribution of simulated transitions (which is identical to the expected percentages across each row in the standard Knox test without covariates).

[Table 7 about here]

Under this definition, for example, Uberlândia's pre-1960 transition was earlier than expected, because 83% of its simulated transitions occur after 1960. In comparison,

[3] The 60% threshold is arbitrary. The higher the threshold, the more unusual the observed transition time must be to qualify as an early or late outlier. We experimented with 50% and 70% cutoffs, and found that results are qualitatively similar to those we report here.

São Paulo's transition during the same period was not early, because its socioeconomic development made it a likely candidate (47% of simulations). The 1960s transitions in Frederico Westphalen and Goiânia were surprisingly early, as was the 1970s transition in Porto Velho (whose relatively low socioeconomic indicators suggested a 63% chance of an even later transition). Manaus's 1980s transition and Balsas's 1990s transition were late, because in both cases local socioeconomic changes predicted earlier events. The other three transitions in the table, like most on the map, were neither early nor late.

[FIGURE 3 about here]

Figure 3 maps the complete set of early and late transitions, thus illustrating some of the sites of transition clusters. By (imperfectly but effectively) filtering out effects of socioeconomic covariates, the map reveals that the surprisingly early transitions to low fertility were clustered in the far South, several bands radiating inland from the capital cities in the states of São Paulo and Paraná, in western Minas Gerais, in Western frontier regions, and in the state of Rio Grande do Norte around its capital, Natal. Late fertility transitions were less clustered, with concentrations mainly at the top edge of the map in the states of Pará (7 of 19 regions late) and Maranhão (11 of 19 regions late).

DISCUSSION

Demographers still debate how well socioeconomic changes can explain the timing of fertility decline, and how much the unexplained variation reflects contagious diffusion of new ideas and behaviors through social interaction. Maps have played an important role in these debates (Coale and Watkins 1986; Watkins 1986), but there has been little use of space-time statistical tests.

The Knox test is a potentially useful tool for demographers seeking to detect epidemic patterns in the frequency or timing of many types of events on maps. In this paper we apply the Knox test to the classic problem, by asking if regional fertility declines in Brazil show evidence of contagion.

Our application demonstrates that demographers may have to adapt, rather than directly borrow, tests from epidemiology and spatial statistics. Brazilian development had a strong spatial pattern throughout our period of study, with large changes in industrialization, infrastructure, and education occurring first in the South and Southeast regions. In effect, one must raise the standard of evidence in any statistical test, to account for spatial patterns and localized changes in socioeconomic variables that could create space-time clustering even in the absence of contagion. Such adaptation will be important when covariates affect demographic events and the risk set changes through time, or when covariates change through time with their own spatial pattern.

We began by recognizing from Figure 1 that fertility transitions clustered in different regions in different periods. The standard Knox test finds very strong evidence of space-time interaction in the timing of these events. We then generalized the Knox test to include covariates, using a proportional hazards model with several likely correlates of fertility decline that are available from census data. Accounting for covariates raises the bar substantially (Figure 2), but not enough to reverse the finding of strong space-time interactions in the observed event rates.

The modified Knox test indicates the presence of strong space-time interaction in our fertility transition data, well beyond that explained by the five covariates in the proportional hazard model. One explanation is that the model is imperfect: it contains

only a partial list of relevant local variables, included covariates are measured with error, and in any case we would not expect a perfect fit to the timing data. Alternatively, the unexpectedly high space-time clustering of fertility transitions may be taken as evidence for social interaction.

In a third step, we developed an additional way to analyze space-time interactions by identifying which regions had unexpectedly early or late fertility transitions, given their covariates. Such analysis is only possible in the modified version of the Knox test, because in the standard test without covariates the timing of decline is distributed identically in all regions (bottom row of Table 7). Figure 1 shows that the early transitions were in the South and Southeast, while late transitions occurred in the North and Northeast. Given the early economic development in the South and Southeast, this space-time interaction sheds little light on whether social interaction or development drove fertility decline. In contrast Figure 3, showing the unexpectedly early and late transitions considering covariates, is much more informative. Here, unexpectedly early transitions are not as concentrated in the South and Southeast, and include a considerable number in the Northeast and Central-West.

Figure 3 generates some new questions about the history of fertility decline in Brazil. What happened in places where fertility declined earlier than would have been expected? Why did transitions begin late in some, but not all, parts of Northern Brazil? One possible explanation for the very earliest transitions (before 1960) in the South is that motivations for having fewer children and knowledge of family planning were ideas that spread at least somewhat independently of development. These ideas may have spread through contact with other cultures and populations— for example, via migrants

from Europe or from neighboring Uruguay. Pioneer regions may then have served as an example to neighboring Southern regions that reached lower fertility levels in the following decade.

The explanation of low fertility as contagion seems plausible with regard to some, but not all, of the earlier than expected transitions in Figure 3. The cluster of early transitions in the Northeastern state of Rio Grande do Norte apparently has a different cause. This state had an early, and reportedly quite effective, effort by a private family planning organization and the state government. The campaign delivered oral contraceptives in over 300 health posts, and gave over 40,000 educational talks about family planning to nearly a million people (Thomé 2006). This effort could well have been responsible for fertility having declined in the 1970s in Rio Grande do Norte, when most of the rest of Northeast only experienced this transition in the following decade.

These interpretations are, of course, speculative. However, they demonstrate the power of looking at the timing of the fertility transition over a large number of relatively small geographic units, and adjusting for covariates as we have done in Figure 3. The modified test directs attention precisely to the space-time clusters that need explanation.

Technical improvements in statistical tests for space-time data will not resolve long-standing debates regarding the nature of fertility transitions. However, we suggest that in this area, as well as in other demographic applications, appropriately adapted formal techniques from epidemiology can be powerful tools for generating interesting questions and highlighting phenomena that might otherwise pass unnoticed.

APPENDIX: EVENT RATES, PERMUTATIONS, AND LIKELIHOODS

Map Likelihoods with Constant Relative Risks

In a space-time point process the likelihood of the n observed events in a map M is (cf Baker 2004, equation 1):

$$L(M) = \left\{\prod_{i=1}^{n} \rho(s_i, t_i)\right\} \exp\left\{-\int_{\mathbf{R}} \rho(s,t)\, ds\, dt\right\} \tag{A1}$$

where ρ represents the intensity or rate of events per space-time volume, conditional on the history of the process through time t. The integral represents appropriate (continuous and/or discrete) summation over the set of space-time combinations with positive risk, denoted **R**. Risk could be zero at some (s,t) combinations – e.g., if location s were not monitored at time t, or if there were a failure at s before time t in a survival model – and any such (s,t) values are excluded from **R**. For a permutation map $M^*(\pi)$ the likelihood becomes

$$L(\pi) = \left\{\prod_{i=1}^{n} \rho(s_{(i)}, t_i)\right\} \exp\left\{-\int_{\mathbf{R}(\pi)} \rho(s,t)\, ds\, dt\right\} \tag{A2}$$

where the notation **R**(π) denotes the possibility that reordering the events could, in principle, change the risk set.

The CRR assumption behind the standard Knox test requires that, under the null hypothesis of no contagion, event rates must be separable into pure time and space effects. Denoting rates under the null hypothesis as ρ_0, the CRR assumption is

$$\rho_0(s,t) = \gamma(t) \cdot \alpha(s) \quad \forall\ s,t \tag{A3}$$

where α(s) is the initial rate in location s at t=0. The common time trend γ(t) depends only on t (and possibly on events occurring prior to t; in this sense it may be a random function). If rates were not separable in this way, then there would be at least one pair of

places for which relative rates change over time. CRR also implies that the risk set cannot change: $\rho_0(A,t)/\rho_0(B,t) = \alpha(A)/\alpha(B)$, so that if locations A and B both have non-zero risks at t=0, they must both have non-zero risks at all times. CRR therefore ensures that reordering event times does not alter the risk set.

Under CRR, risk sets **R** and local rates ρ must be identical for observed and reordered maps, making the integral terms identical in Equations (A1) and (A2). The likelihood ratio under the null hypothesis is therefore

$$\text{(A4)} \qquad \frac{L_0(\pi)}{L_0(M)} = \frac{\prod_i \rho_0(s_{(i)},t_1)}{\prod_i \rho_0(s_i,t_1)} = \frac{\prod_i \gamma(t_i)\alpha[s_{(i)}]}{\prod_i \gamma(t_i)\alpha[s_i]} = \frac{\prod_i \alpha[s_{(i)}]}{\prod_i \alpha[s_i]} = 1$$

where the last equality occurs because permutation alters the order of multiplication in the numerator and denominator, but not the product. Equation (A4) thus shows that under CRR all n! possible permutation maps (including the original map, π=[1,2,…,n]) must have identical likelihood $L_0(\pi) = 1/n!$.

Map Likelihoods with Covariates

A useful extension for including covariates in a non-contagious process is

$$\text{(A5)} \qquad \rho_0(s,t) = \gamma(t)\ \exp\left[\ \beta'\ z(s,t)\ \right]$$

where $\gamma(t)$ is a baseline time trend, β is a parameter vector, and z is a vector of covariates (cf. Diggle 1990, equation 14). As in a typical regression model, inference is conditional on observed covariate values, so that $z(s,t)$ is a non-random function. In a model with covariates, lack of contagion is no longer synonymous with CRR, because variations in $z(s,t)$ can cause space-time interactions in rates.

This formulation intentionally mimics a proportional hazards model in which covariates affect event risks (Cox 1972). However, it is not necessary to assume a survival process in which events remove units (in this case, locations) from the risk set. We will use this flexibility in several ways in the discussion that follows.

Combining equations (A5) and (A1) and factoring out common time trends yields an expression for the likelihood of a permutation map with covariates under the null hypothesis of no contagion:

$$\text{(A6)} \qquad L_0(\pi \mid \beta) \propto \exp\left\{\beta' \sum_{i=1}^{n} z(s_{(i)}, t_i)\right\} \exp\left\{-\int_{\mathbf{R}(\pi)} \gamma(t) \exp[\beta' z(s,t)]\, ds\, dt\right\}$$

Null Distribution of the Knox Statistic across Permutation Maps

A permutation map's Knox statistic is $X^*(\pi)=X[M^*(\pi)]$, so the *p*-value associated with an observed Knox statistic X is

$$\text{(A7)} \qquad p_0(X) = \sum_{k=1}^{n!} L_0(\pi_k \mid \beta) \cdot I\left[X^*(\pi_k) \geq X\right]$$

With covariates, the likelihoods L_0 are given in Equation (A6). In general L_0 depends on the particular ordering $\pi=(1)\ldots(n)$, and expression for the null distribution in Equation (A7) cannot be further simplified. There are three special cases of interest, however, in which the distribution is more tractable. We discuss these in turn.

Case 1: Constant Relative Risks (β irrelevant)

The first special case is that assumed by Knox (1964) in which local factors, such as population, change over the monitoring period in ways that leave relative risks

constant. This is the CRR assumption, and if it holds (Equation A4) then all permutation likelihoods are identical.

In a model with covariates CRR requires extremely strong assumptions. Specifically, CRR holds only if (1) the risk set is unaltered by changing the order of events, and (2) covariates change so that for all permutations the likelihood ratio

$$\text{(A8)} \qquad \frac{L_0(\pi \mid \beta)}{L_0(M \mid \beta)} = \exp\left\{\beta' \sum_{i=1}^{n}\left[z(s_{(i)}, t_i) - z(s_i, t_i)\right]\right\} = 1$$

This ratio equals unity for all permutations π if covariates are irrelevant ($\boldsymbol{\beta}=\mathbf{0}$), or if covariates change in very restrictive ways that leave the exponential term equal to zero under all permutations.

Case 2: Constant Risk Sets and Shifting Populations (β known)

Kulldorff and Hjalmars (1999) explore a second special case that is important in demographic applications spanning long periods of time. They note that temporal shifts in the distribution of population across space could bias the standard Knox test in favor of finding non-existent epidemicity. Their proposed alternative test is complex, mainly because it does not condition on the marginal distributions of observed locations $\{s_1 \ldots s_n\}$ and times $\{t_1 \ldots t_n\}$.

Our framework suggests an alternative solution to the problem of shifting population that is simple and arguably more consistent with the standard Knox Monte Carlo method. Specifically, the Kulldorff and Hjalmars (1999) problem can be expressed in a proportional hazards framework as a special case in which risk sets are unchanging, *z*

is a scalar equal to the logarithm of population, $z(s,t)=\ln N(s,t)$, and $\beta=1$. In this case,

$$L_0(\pi|\beta) \propto \prod_{i=1}^{n} N\left[s_{(i)}, t_i\right] = N^*(\pi) \tag{A9}$$

In other words, a permutation of observed times and places is more likely if it allocates events to high-population cells. This immediately suggests a modified Monte Carlo strategy for the Knox test over K permutations drawn uniformly from the set of $n!$ possibilities:

$$\hat{p}_0(X) = \left(\sum_{k=1}^{K} N^*(\pi_k) \cdot I\left[X^*(\pi_k) > X\right]\right) \Big/ \left(\sum_{k=1}^{K} N^*(\pi_k)\right) \tag{A10}$$

We will pursue this approach in another paper; we merely note here that it is another potentially interesting extension of the Knox method for demographic and epidemiological studies.

Case 3: Discrete Locations and Survival Processes with Covariates (β estimable)

The third special case, which is the main focus of this paper, occurs with a survival process over discrete spatial areas. In such a model, each area has a single transition or failure time, after which it is removed from the risk set. The permutation π in this case determines the order of transitions/failures, which is precisely the information used to estimate Cox's (1972) proportional hazard model. In particular, in the model set out here, the likelihood of failures in order (1)…(n) under the null hypothesis of no contagion is (Cox 1972, equation 12; Kalbfleisch and Prentice 1973, equation 2):

$$L_0(\pi|\beta) = \prod_{i=1}^{n} \left\{ \exp\left[\beta' z(s_{(i)}, t_i)\right] \Big/ \sum_{j \in R_{(i)}} \exp\left[\beta' z(s_j, t_i)\right] \right\} \tag{A11}$$

where $R_{(i)}$ represents the set of areas $\{(i)\ldots(n)\}$ with failures on or after time t_i. Permutation likelihoods in this case depend on covariate effects β, and on the pattern of covariates over time and space.

Our modification of the Knox test uses a Monte Carlo simulation in which map orderings that are more likely, given the observed histories of relevant covariates, are more likely to be sampled. Unequal-probability sampling of permutations is easy with the *sample* function in the *R* statistical language (Ihaka and Gentleman 1996), using probabilities $\exp\left(\hat{\beta}'\mathbf{z}\right)$ as arguments. With sampling of permutations proportional to their probabilities, the simple average

$$\hat{p}_0(X) = \frac{1}{K}\sum_{k=1}^{K} I\left[X^*(\pi_k) \geq X\right] \tag{A12}$$

converges to $p_0(X) = \Pr[X^* \geq X \mid H_0]$ as K increases.

Modified Monte Carlo procedures with covariate effects and unequal permutation probabilities are a key contribution of this paper. Standard Knox approaches, based either on analytical approximations or on simulations of the null distribution of X, are biased when covariates and risk sets change over the monitoring period. Such conditions are likely in demographic studies, and our modified procedures eliminate this source of bias.

REFERENCES

Anselin, L. 1988. *Spatial Econometrics*. Boston: Kluwer.

Baker, R.D. 2004. "Identifying space-time disease clusters". *Acta Tropica* 91:291-299.

Baller, R.D. and K. Richardson. 2002. "Social Integration, Imitation, and the Geographic Patterning of Suicide". *American Sociological Review* 67:872-888.

Baller, R.D., L. Anselin, G. Deane, S. Messner and D. Hawkins. 2001. "Structural Covariates of U.S. County Homicide Rates: Incorporating Spatial Effects". *Criminology* 39:561-590.

Balabdaoui, F., J.P. Bocquet-Appel , C. Lajaunie, and S.I. Rajan. 2001. "Space-time evolution of fertility transition in India (1961-91)". *International Journal of Population Geography*, 7(2):129-148.

Barton, D.E. and F.N. David. 1966. "The random intersection of two graphs", pp. 445-459 in *Research Papers in Statistics: Festschrift for J. Neyman*, edited by F.N. David. Wiley: New York.

Bhopal, R.S., P.J. Diggle, and B.S. Rowlingson. 1992. "Pinpointing clusters of apparently sporadic Legionnaire's disease." *British Medical Journal* 304:1022-1027.

Birch, J.M., F.E. Alexander, V. Blair, O.B. Eden, G.M. Taylor, and R.J.Q. McNally. 2000. "Space-time clustering patterns in childhood leukaemia support a role for infection". *British Journal of Cancer* 82:1571-1578.

Bocquet-Appel, J.P. and L. Jakobi. 1998. "Evidence for a spatial diffusion of contraception at the onset of the fertility transition in Victorian Britain". *Population: An English selection, Special issue New advances in Social Sciences* 10(1):181-204.

Bocquet-Appel, J.P., S.I. Rajan, J.N. Bacro, and C. Lajaune. 2002. "The onset of India's fertility transition". *European Journal of Population* 18:211-232.

Coale, A.J. 1973. "The Demographic Transition Reconsidered". In International Union for the Scientific Study of Population, International Population Conference, Liège, Belgium. Vol. 1:53-72.

Coale, A.J. and R. Treadway. 1986. "A summary of the changing distribution of overall fertility, marital fertility, and the proportion married in the provinces of Europe", pp. 31-181 in *The decline of fertility in Europe*, edited by A.J. Coale and S.C. Watkins. Princeton, New Jersey: Princeton University Press.

Coale, A.J. and S.C. Watkins (eds.). 1986. *The decline of fertility in Europe*. Princeton, NJ: Princeton University Press.

Cox, D.R. 1972. "Regression Models and Life-Tables". *Journal of the Royal Statistical Society B* 34(2):187-220.

Diggle, P.J. 1990. "A Point Process Modelling Approach to Raised Incidence of a Rare Phenomenon in the Vicinity of a Prespecified Point". *Journal of the Royal Statistical Society A* 153(3):349-362.

Guinnane, T.W., B.S. Okun, and J. Trussell. 1994. "What Do We Know About the Timing of Fertility Transitions in Europe?". *Demography* 31(1), 1-20.

Ihaka, R. and R. Gentleman, 1996. "R: A language for data analysis and graphics". *Journal of Computational and Graphical Statistics* 5(3):299-314.

Kalbfleisch, J.D. and R.L. Prentice. 1973. "Marginal Likelihoods based on Cox's Regression and Life Model". *Biometrika* 60(2):267-278.

Knox, E.G. 1964a. “The Detection of Space-Time Interactions”. *Applied Statistics* 13(1):25-30.

--------- . 1964b. “Epidemiology of Childhood Leukemia in Northumberland and Durham”. *British Journal of Preventative and Social Medicine* 18(1):17-24.

---------. 2002. “An epidemic pattern of murder”. *Journal of Public Health Medicine* 24:34-37.

Land, K. and G. Deane, 1992. “On the Large-Sample Estimation of Regression Models with Spatial- or Network-Effects Terms: A Two-State Least Squares Approach”, in *Sociological Methodology*, edited by P.V. Marsden. Washington: American Sociological Association.

Land, K., G. Deane and J. Blau. 1991. “Religious Pluralism and Church Membership: A Spatial Diffusion Model.” *American Sociological Review* 56:237-49.

Kulldorff, M. and U. Hjalmars, 1999. “The Knox Method and Other Tests for Space-Time Interaction”. *Biometrics* 55(2):544-552.

Machado-Coelho, G.L., R. Assunção, W. Mayrink and W.T. Caiaffa. 1999. “American Cutaneous Leishmaniasis in Southeast Brazil: Space-time Clustering”. *International Journal of Epidemiology* 28:982-989.

McKenzie, N., S. Landau, N. Kapur, J. Meehan, J. Robinson, H. Bickley, R. Parsons and L. Appleby. 2005. "Clustering of suicides among people with mental illness". *British Journal of Psychiatry* 187: 476-480

McNally, R.J.Q., D.P. Cairns, O.B. Eden, F.E. Alexander, G.M. Taylor, A.M. Kelsey and J.M. Birch, 2002. “An Infectious Aetiology for Childhood Brain Tumours?

Evidence from Space-Time Clustering and Seasonality Analyses". *British Journal of Cancer* 86:1070-1077.

Montgomery, M.R. and J.B. Casterline, 1996. "Social Learning, Social Influence, and New Models of Fertility". *Population and Development Review 22, Supplement: Fertility in the United States: New Patterns, New Theories*, pp. 151-175.

Morris, J.K., E. Alberman and D. Mutton. 1998. "Is there evidence of clustering in Down syndrome?". *International Journal of Epidemiology* 27:495-498.

Morrison, A.C., A. Getis, M. Santiago, J.G. Rigau-Perez and P. Reiter. 1998. "Exploratory Space-Time Analysis of Reported Dengue Cases during an Outbreak in Florida, Puerto Rico, 1991-1992". *American Journal of Tropical Medicine and Hygiene* 58(3):287-298.

Potter, J.E., C.P. Schmertmann and S.M. Cavenaghi, 2002. "Fertility and Development: Evidence from Brazil". *Demography* 39(4):739-761.

Samuelsson, U., C. Johansson, J. Carstensen and J. Ludvigsson. 1994. "Space-time clustering in insulin-dependent diabetes mellitus (IDDM) in South-East Sweden". *International Journal of Epidemiology* 23:138-142.

Thomé, M. 2006. Personal communication containing program statistics compiled by BEMFAM (Bem-Estar Familiar no Brasil).

Tobin, P.C. 2007. "Space-time patterns during the establishment of a nonindigenous species". *Population Ecology* 49(3):257-263.

Tolnay, S.E. 1995. "The Spatial Diffusion of Fertility: A Cross-Sectional Analysis of Counties in the American South, 1940". *American Sociological Review* 60(2):299-308.

Tolnay, S.E., G. Deane and E.M. Beck, 1996. "Vicarious Violence: Spatial Effects on Southern Lynchings, 1890-1919". *American Journal of Sociology* 102(3):788-815.

Watkins, S.C. 1986. "Conclusions", pp. 420-449 in *The decline of fertility in Europe*, edited by A.J. Coale and S.C. Watkins. Princeton, NJ: Princeton University Press.

Figure 1: Fertility Transitions by Decade, Brazil 1960-2000. Shaded regions had TFR < 4.5 by the end of the period, or had at least a 20% decline over the period.

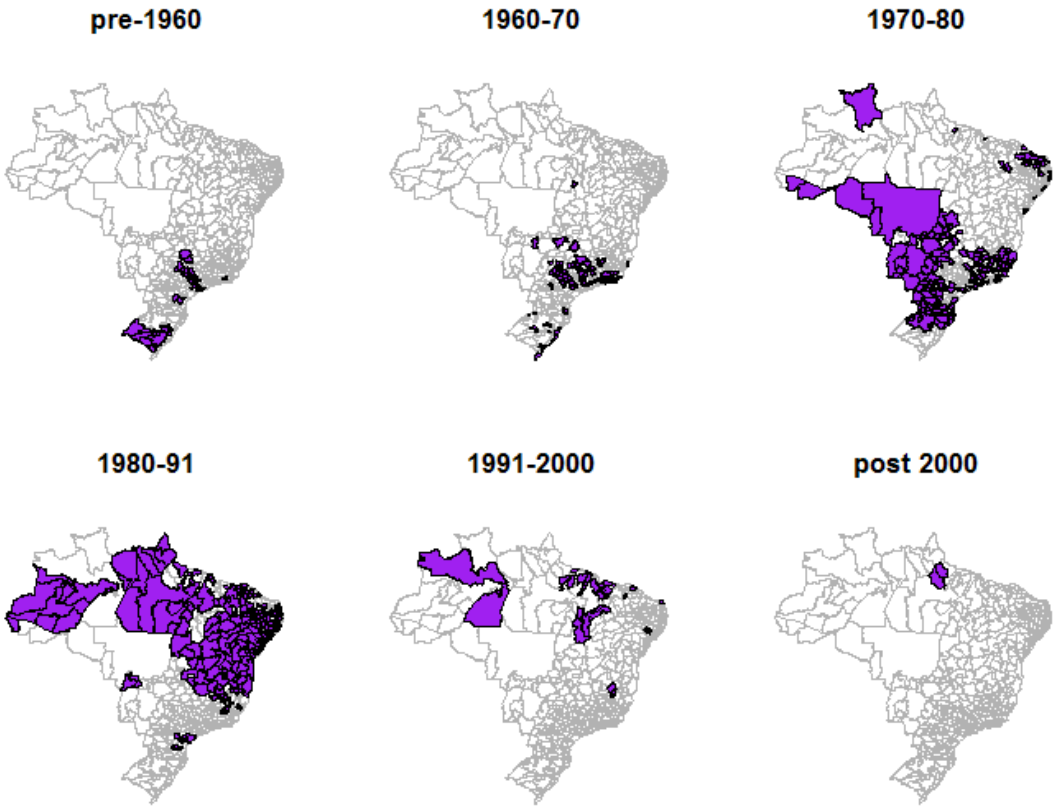

Figure 2: Standard Knox test, and the modified test with covariate effects. 10000 Monte Carlo permutations. Observed X=1119. Median X*=545 in standard test, 705 in modified test. 5% significance level at X=578 in standard test, 753 in modified test.

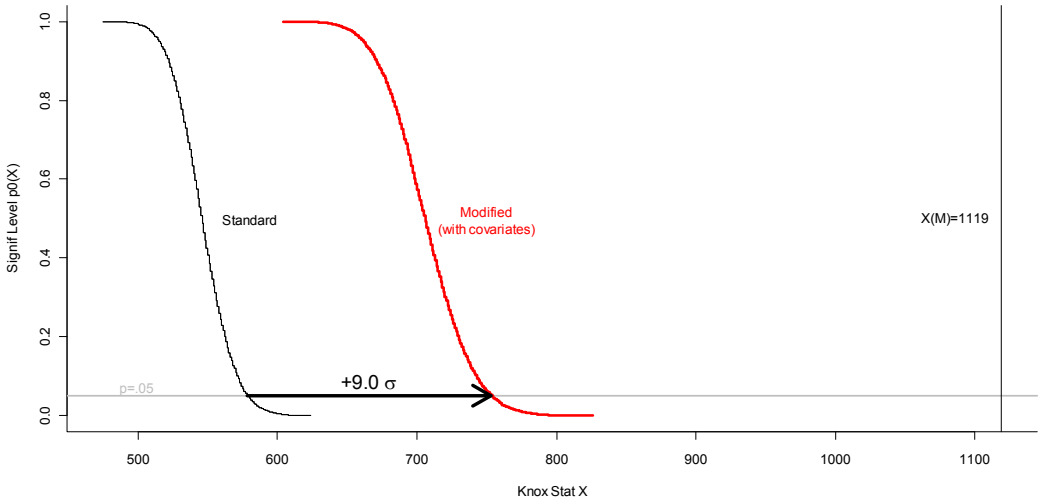

Figure 3: Early (light grey) and Late (dark) Regional Fertility Transitions relative to Socioeconomic Conditions, Brazil 1960-2000. Early [late] transitions are those for which at least 60% of simulated transition times in the modified Knox test occurred after [before] the observed transition.

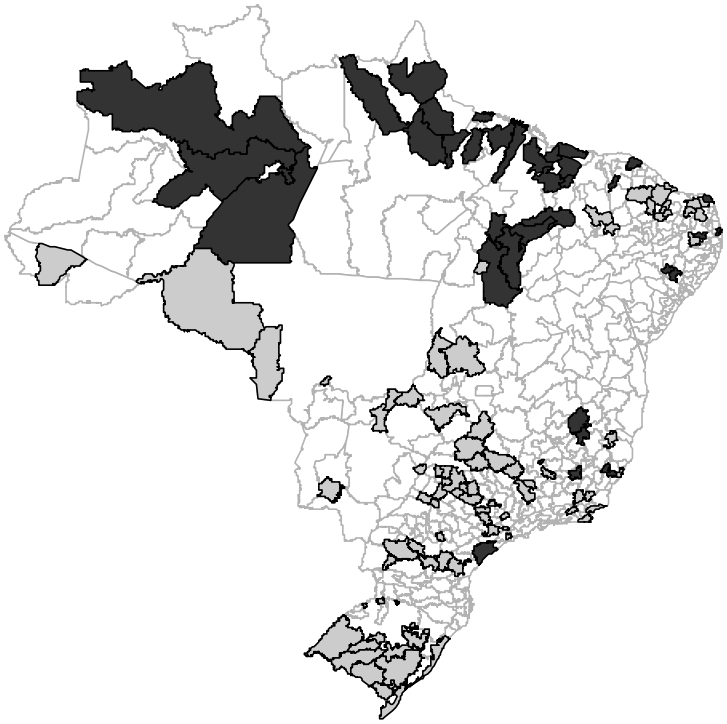

Table 1: Distribution of fertility transitions by intercensal period

	Pre 1960	1960-70	1970-80	1980-91	1991-00	Post 2000
# Regions at Risk	395[1]	470	405	227	29	2
# Transitions	32	65	178	198	27	2
Period failure rate (%) [2]	8	14	44	87	93	100
Fraction of all transitions (%)	6	13	35	39	5	0.4

Notes:
1. 1960 data are unavailable for 107 of the 502 regions.
2. Period failure rate = # Transitions / # Regions at Risk

Table 2: Estimated period TFR and inferred transition times for selected regions

	Period TFR					
Region	*1960*	*1970*	*1980*	*1991*	*2000*	*Transition*
Porto Velho, RO	--	7.0	5.3*	3.4^{X*}	2.5^{X*}	1970-1980
Manaus, PA	--	6.0	5.2	3.2^{X*}	2.6^{X}	1980-1991
Santarém, PA	--	7.1	6.6	4.4^{X*}	3.3^{X*}	1980-1991
Balsas, MA	--	6.0	6.0	5.1	3.6^{X*}	1991-2000
João Pessoa, PB	5.3	4.9	3.9^{X*}	2.5^{X*}	1.9^{X*}	1970-1980
Uberlândia, MG	4.5^{X}	3.8^{X}	2.9^{X*}	2.2^{X*}	1.8^{X}	Pre 1960
São Paulo, SP	3.1^{X}	3.0^{X}	2.8^{X}	2.1^{X*}	1.9^{X}	Pre 1960
Frederico Westphalen, RS	7.6	5.7*	3.9^{X*}	2.9^{X*}	2.2^{X*}	1960-1970
Goiânia, GO	4.9	4.2^{X}	3.1^{X*}	2.1^{X*}	1.8^{X}	1960-1970
BRAZIL AVERAGE	6.0	5.8	4.9	3.4	2.6	

Notes:
1. ^{X} indicates TFR < 4.5 (transition condition 1); * indicates an intercensal decline of >20% (condition 2). In our definition a transition has occurred if either condition holds. Post-transition cells are shaded.
2. BRAZIL AVERAGE is the simple mean across all 502 regions, not population-weighted.
3. 1960 demographic census data are unavailable from any source for 107 regions, primarily in the North. The 1960 average TFR on the bottom row excludes these 107 regions. For these regions we assume the earliest possible mortality transition was in 1960-70.

Table 3: Knox statistic for the fertility transition map

	Event pairs (*i*,*j*)	
Close in Space \ *Close in Time*	*NO (τ_{ij} = 0)*	*YES (τ_{ij} = 1)*
NO (σ_{ij} = 0)	86889	37065
YES (σ_{ij} = 1)	678	***X* = 1119**

Table 4: Proportional Hazard Regression Estimates for Fertility Transition

	$\hat{\beta}$	exp($\hat{\beta}$)	se($\hat{\beta}$)	p
Percent of Households with Electricity	.039	1.039	.004	.000
Percent of Labor Force in Primary Sector	.016	1.017	.004	.000
Percent Female Labor Force Participation	-.018	0.983	.009	.047
Average Years of Female Education	.456	1.577	.085	.000
Percent Catholic	-.013	0.987	.007	.075

Table 5. Standard and Modified Knox Tests with Alternative Definitions of Fertility Transition, Spatial Proximity, and Temporal Proximity. Rightmost column shows the shift in the 95%ile of the X* distribution, measured in standard deviations, when using the modified (with covariate) test.

		Proximity					Standard X* distribution (without covariates)			Modified (with covariates)	
Case[2]	Transition[3]	Space[4]	Time[5]	X	N_s	N_t	σ	μ	95%ile	95%ile	Shift / σ
1	<4.5 | >20%	100km | 1lag	same pd	1119	1797	38184	19.4	545	578	753	+9.0
2	<4.5 | >20%	100km	same pd	720	1184	38184	17.5	359	388	502	+6.5
3	<4.5 | >20%	1lag	same pd	878	1404	38184	17.2	426	456	593	+8.0
4	<4.5 | >20%	250km | 2lag	same pd	4116	7677	38184	44.2	2331	2403	3079	+15.3
5	<4.5 | >20%	250km	same pd	3433	6415	38184	52.7	1947	2039	2577	+10.2
6	<4.5 | >20%	2lag	same pd	2703	4770	38184	33.4	1447	1501	1962	+13.8
7	<4.5 | >20%	100km | 1lag	same | adj pd	1720	1797	92478	19.3	1322	1352	1507	+8.0
8	<4.5 | >20%	100km	same | adj pd	1141	1184	92478	20.1	870	902	998	+4.8
9	<4.5 | >20%	1lag	same | adj pd	1343	1404	92478	15.8	1033	1059	1181	+7.7
10	<4.5 | >20%	250km | 2lag	same | adj pd	7099	7677	92478	55.9	5645	5739	6317	+10.3
11	<4.5 | >20%	250km	same | adj pd	5967	6415	92478	71.0	4716	4840	5291	+6.4
12	<4.5 | >20%	2lag	same | adj pd	4445	4770	92478	35.0	3505	3563	3969	+11.6
13	<4.0	100km | 1lag	same pd	1017	1797	32138	18.6	459	489	717	+12.3
14	<4.0	100km	same pd	668	1184	32138	15.8	303	329	482	+9.7
15	<4.0	1lag	same pd	824	1404	32138	16.4	359	385	570	+11.3
16	<4.0	250km | 2lag	same pd	3737	7677	32138	42.5	1963	2030	2909	+20.7
17	<4.0	250km	same pd	3159	6415	32138	42.8	1639	1711	2467	+17.7
18	<4.0	2lag	same pd	2492	4770	32138	31.1	1221	1273	1859	+18.8
19	<4.0	100km | 1lag	same | adj pd	1703	1797	84090	20.4	1202	1236	1512	+13.5
20	<4.0	100km	same | adj pd	1121	1184	84090	20.5	791	826	1007	+8.8
21	<4.0	1lag	same | adj pd	1342	1404	84090	17.5	940	969	1194	+12.9
22	<4.0	250km | 2lag	same | adj pd	7117	7677	84090	56.8	5133	5230	6378	+20.2
23	<4.0	250km	same | adj pd	5973	6415	84090	68.5	4289	4399	5356	+14.0
24	<4.0	2lag	same | adj pd	4476	4770	84090	36.5	3190	3248	4002	+20.7
25	>10%	100km | 1lag	same pd	1070	1797	42808	20.2	611	645	759	+5.6
26	>10%	100km	same pd	694	1184	42808	18.4	404	434	507	+4.0
27	>10%	1lag	same pd	880	1404	42808	17.2	478	505	597	+5.3
28	>10%	250km | 2lag	same pd	4256	7677	42808	45.9	2611	2690	3142	+9.8
29	>10%	250km	same pd	3592	6415	42808	53.1	2186	2280	2663	+7.2
30	>10%	2lag	same pd	2732	4770	42808	33.4	1625	1680	1976	+8.9
31	>10%	100km | 1lag	same | adj pd	1751	1797	103805	15.9	1483	1510	1614	+6.5
32	>10%	100km	same | adj pd	1150	1184	103805	16.8	977	1006	1073	+4.0
33	>10%	1lag	same | adj pd	1373	1404	103805	13.7	1159	1181	1263	+6.0
34	>10%	250km | 2lag	same | adj pd	7411	7677	103805	43.7	6335	6407	6862	+10.4
35	>10%	250km	same | adj pd	6186	6415	103805	55.8	5294	5384	5755	+6.6
36	>10%	2lag	same | adj pd	4626	4770	103805	29.8	3938	3987	4261	+9.2
AVERAGE OVER ALL 36 EXPERIMENTS											+10.4

Notes:

1. Vertical bars denote “or” conditions
2. Case 1 results are from 10000 random maps, as described in the text. Cases 2-36 are from 1000 random maps.
3. Transition definitions are in terms TFR at end of intercensal period, or % decline in TFR over period
4. Spatial proximity between transitions in terms of distance between centroids, or number of spatial lags between regions
5. Temporal proximity between transitions defined as “happened in same intercensal period”, or “same or adjacent periods”

Table 6: Distribution of Early and Late Transitions, by intercensal period

	Observed Transition Time						
	Pre 1960	1960-70	1970-80	1980-91	1991-00	Post 2000	ALL
Early	24	31	31	0	0	0	86
Neither Early nor Late	8	34	147	189	1	0	379
Late	0	0	0	9	26	2	37
Total	32	65	178	198	27	2	502

Table 7: Simulated and observed fertility transition times for example regions. Observed transition times are shaded. Early [late] transitions are those for which at least 60% of simulated transition times in the modified Knox test occurred after [before] the observed transition.

	Simulated Transition Times (% of samples)						
Region	Pre 1960	1960-70	1970-80	1980-91	1991-00	Post 2000	Early/Late?
Porto Velho, RO	0	8	29	62	1	0	Early
Manaus, PA	0	20	66	15	0	0	Late
Santarém, PA	0	6	41	52	1	0	--
Balsas, MA	0	5	13	71	11	0	Late
João Pessoa, PB	11	22	61	6	0	0	--
Uberlândia, MG	17	32	47	3	0	0	Early
São Paulo, SP	47	32	20	0	0	0	--
Frederico Westphalen, RS	5	11	50	34	0	0	Early
Goiânia, GO	13	27	58	2	0	0	Early
ALL REGIONS[1]	*6*	*13*	*35*	*39*	*5*	*0*	--

Note:

1. From Table 1. In a standard Knox test without covariates, every region would have this distribution of simulated transition times.