Social Science & Medicine 91 (2013) 58–66

Contents lists available at SciVerse ScienceDirect

Social Science & Medicine

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

Investigating the relationship between neighborhood poverty and mortality risk: A marginal structural modeling approach

D. Phuong Do [a,d,*], Lu Wang [b], Michael R. Elliott [b,c]

[a] Department of Health Services Policy and Management, University of South Carolina, 800 Sumter St, Columbia, SC 29208, USA
[b] Department of Biostatistics, University of Michigan, USA
[c] Survey Methodology Program, Institute for Social Research, University of Michigan, USA
[d] Zilber School of Public Health, University of Wisconsin, Milwaukee, 1240 N. 10th Street, Milwaukee, WI 53201, USA

ARTICLE INFO

Article history:
Available online 14 March 2013

Keywords:
Marginal structural modeling
Neighborhood effects
Longitudinal analysis
Causal modeling
Sensitivity analysis

ABSTRACT

Extant observational studies generally support the existence of a link between neighborhood context and health. However, estimating the causal impact of neighborhood effects from observational data has proven to be a challenge. Omission of relevant factors may lead to overestimating the effects of neighborhoods on health while inclusion of time-varying confounders that may also be mediators (e.g., income, labor force status) may lead to underestimation. Using longitudinal data from the 1990 to 2007 years of the Panel Study of Income Dynamics, this study investigates the link between neighborhood poverty and overall mortality risk. A marginal structural modeling strategy is employed to appropriately adjust for simultaneous mediating and confounding factors. To address the issue of possible upward bias from the omission of key variables, sensitivity analysis to assess the robustness of results against unobserved confounding is conducted. We examine two continuous measures of neighborhood poverty – single-point and a running average. Both were specified as piece-wise linear splines with a knot at 20 percent. We found no evidence from the traditional naïve strategy that neighborhood context influences mortality risk. In contrast, for both the single-point and running average neighborhood poverty specifications, the marginal structural model estimates indicated a statistically significant increase in mortality risk with increasing neighborhood poverty above the 20 percent threshold. For example, below 20 percent neighborhood poverty, no association was found. However, after the 20 percent poverty threshold is reached, each 10 percentage point increase in running average neighborhood poverty was found to increase the odds for mortality by 89 percent [95% CI = 1.22, 2.91]. Sensitivity analysis indicated that estimates were moderately robust to omitted variable bias.

Published by Elsevier Ltd.

Introduction

Extant observational studies generally support the existence of a link between neighborhood context and health (Diez Roux & Mair, 2010; Kawachi & Berkman, 2003; Pickett & Pearl, 2001; Robert, 1999; Yen & Syme, 1999). Most of the existing literature has focused on neighborhood demographic or socioeconomic characteristics, particularly neighborhood poverty and disadvantage (Meijer, Rohl, Bloomfield, & Grittner, 2012; Riva, Gauvin, & Barnett, 2007; Robert, 1999). Other research has examined the social aspects of neighborhood environments (e.g., social capital, trust, and crime) in relation to health (Sampson, Morenoff, & Gannon-Rowley, 2002). More recently, attention has been paid to the built environment, such as housing conditions, ambient air quality, and urban form (Frank, Engelke, & Schmid, 2003). The various health outcomes that have been linked to neighborhood context include, among others, mortality (Leclere, Rogers, & Peters, 1998; Sloggett & Joshi, 1998; Yen & Syme, 1999), infectious disease (Acevedo-Garcia, 2000, 2001), low birthweight (Morenoff, 2003; O'Campo, Xue, Wang, & Caughy, 1997; Roberts, 1997; Sastry & Hussey, 2003) cigarette smoking (Diez Roux, Merkin, Hannan, Jacobs, & Kiefe, 2003; Duncan, Jones, & Moon, 1999; Immo Kleinschmidt, 1995) and diet (Morland, Wing, & Roux, 2002). These associations between health and place generally remain statistically significant even after adjusting for various individual-level socioeconomic characteristics (c.f., Oreopoulos, 2003; Reijneveld & Schene, 1998; Andrew Sloggett & Joshi, 1994).

However, estimating the causal impact of neighborhood effects has proven to be a challenge. Since random assignment into different neighborhood context is cost-prohibitive and infeasible on a large scale, most neighborhood-effect studies have relied on observational data. And as with inferences from all observational

* Corresponding author. Zilber School of Public Health, University of Wisconsin, Milwaukee, 1240 N. 10th Street, Milwaukee, WI 53201, USA. Tel.: +1 414 227 3001.
E-mail address: dphuong@uwm.edu (D.P. Do).

0277-9536/$ – see front matter Published by Elsevier Ltd.
http://dx.doi.org/10.1016/j.socscimed.2013.03.003

studies, a constant concern with neighborhood effect estimates is that they are biased due to unobserved confounding. For example, if persons who are inherently less inclined to live a healthy lifestyle were also less likely to seek neighborhoods with the physical and social environment that supports physical activity and well-being, attributes that are negatively correlated to neighborhood disadvantage, then not accounting for this factor would lead to spurious neighborhood findings. Thus, estimates of significant neighborhood effects may be overestimating the true impact.

However, estimating neighborhood effects presents additional challenges because neighborhood exposure is dynamic. To the extent that neighborhood context varies across time, single-point measures which do not allow for any lagged or long-term effects of neighborhoods on health may also bias estimates. Cross-sectional analyses implicitly assume that either 1) neighborhood context throughout the life-course is relatively constant or 2) neighborhood context has an immediate impact on the health outcome(s) being investigated. While cross-sectional neighborhood effect estimates are rarely explicitly interpreted as either a short-term or long-term exposure, evidence suggests that single-point estimates underestimate the long-term exposure to neighborhood context (Do, 2009). Additionally, an important consideration, which has only recently gained attention, is that factors such as income and employment status may themselves be affected by prior neighborhood conditions and are consequently confounders and mediators of neighborhood-health effects simultaneously. For example, exposure to poor socioeconomic environments during childhood and adolescence has been shown to detrimentally impact educational outcomes (Harding, 2003; Sampson, Sharkey, & Raudenbush, 2008; Wodtke, Harding, & Elwert, 2011). Given the strong linkage between educational attainment and adult health (Adler, Boyce, Chesney, Folkman, & Syme, 1993; Cutler & Lleras-Muney, 2010; Kimbro, Bzostek, Goldman, & Rodriguez, 2008), adjusting for educational attainment prohibits recovering the indirect effect of neighborhoods on health via education, resulting in biased estimates of the total effect. Other similarly important potential mediators include marital status, income, and employment status, which, when collectively controlled for in neighborhood-health models, may result in non-trivial bias of neighborhood effects. Because these factors are in the causal pathway, a major concern in adjusting for these characteristics is the possible underestimation of the total effect. While it may be intuitive to interpret these estimates as the direct impact of neighborhoods that do not pass through these mediators, this is not the case. Models that include time-varying covariates that are simultaneous mediators and confounders produce biased estimates of the direct impact (Nandi, Glymour, Kawachi, & VanderWeele, 2012; Robins, 1999).

Hence, researchers have longed faced a conundrum in which omission of these factors leads to upwardly biased effects while inclusion likely leads to downwardly biased estimates. To avoid the criticism of egregiously overestimating the effects of neighborhoods, studies have taken the conservative approach and conventionally included these simultaneous mediators and confounders in regression adjustments. However, with longitudinal data, appropriate adjustments for the effects of time-dependent confounders (e.g., income) to recover causal estimates of time-dependent treatments (e.g., neighborhood poverty) can be achieved by using marginal structural modeling (MSM) (Robins, Hernan, & Brumback, 2000).

While the usage of MSM strategies is increasing in epidemiologic research (e.g., Bodnar, Davidian, Siega-Riz, & Tsiatis, 2004; Cook, Cole, & Hennekens, 2002), published work that has applied MSM strategies to investigate the effects of neighborhoods is still sparse. The few studies that have examined neighborhood effects on health outcomes via an MSM approach have found exposure to neighborhood disadvantage to increase the level of alcohol consumption and propensity for binge drinking, self-rated poorer health and decrease injection cessation among drug users (Cerda, Diez-Roux, Tchetgen, Gordon-Larsen, & Kiefef, 2010; Glymour, Mujahid, Wu, White, & Tchetgen Tchetgen, 2010; Nandi et al., 2010). Within the broader scope of neighborhood-effect studies in general, several studies have used MSM to examine the effects of childhood exposure to neighborhood deprivation on educational outcomes (Harding, 2003; Sampson et al., 2008; Sharkey & Elwert, 2011; Wodtke et al., 2011). These studies found childhood exposure to neighborhood deprivation to adversely affect high school graduation rates, cognitive ability, and verbal ability.

This paper seeks to address the above-mentioned issues that may bias neighborhood effect estimates in both directions. Using longitudinal data, we address the issue of possible downward bias due to adjusting for mediating factors by employing a marginal structural modeling strategy to investigate the link between neighborhood poverty and mortality risk. It then compares conventional naïve estimates to those recovered from marginal structural modeling. To address the issue of possible upward bias due to omitted variables, we then conduct a sensitivity analysis to assess the robustness of results against unobserved confounding.

Methods

Estimating treatment effects from MSMs is a two-stage process. In stage 1, the treatment assignment (here neighborhood poverty) is modeled. Since we treat neighborhood poverty as a continuous measure, we must estimate each person's probability density of having his/her own treatment at each time point, rather than as probabilities associated with a binary variable. The probability densities are then used to derive stabilized inverse probability of treatment (IPT) weights for each person (Robins et al., 2000). The higher the probability density of an individual to receive the observed treatment, the lower the weight the individual is assigned. Conversely, the lower the probability density that the individual received the observed treatment (based on observed covariates), the larger the weight the individual is assigned. Intuitively, this can be interpreted as upweighting individuals whose neighborhood exposure is underrepresented compared to what would have been observed through random assignment, and downweighting individuals whose neighborhood exposure is overrepresented. This weighting scheme transforms the covariate distributions of the sample population so that the weighted covariate distributions of the group across treatment levels become comparable. Essentially, the process creates a pseudopopulation in which treatment and the time-dependent confounders are no longer associated.

In stage 2, the causal parameter in this pseudopopulation is recovered by estimating a weighted regression model (observed population weighted by the stabilized IPT weight) on the outcome of interest. Under conditions of conditional exchangeability (i.e., treated assignment is effectively randomized, conditional on past treatment history and confounder history), this process will yield an unbiased estimate of the causal parameter of interest. Thus, by accounting for confounding without including these confounders in the structural part of the model, the IPT weight avoids the problems of biased estimation that arise when time-dependent confounders are inappropriately adjusted for via stratification or traditional regression approaches. This allows separating the time-dependent covariates confounder adjustment – which we want to incorporate in our modeling approach – from the mediation adjustment, which we want to avoid when assessing the total impact (direct and indirect) of the treatment on the outcome (Joffe, Ten Have, Feldman, & Kimmel, 2004; Robins et al., 2000).