CAUSAL INFERENCE WITH SPATIO-TEMPORAL DATA:

ESTIMATING THE EFFECTS OF AIRSTRIKES ON INSURGENT VIOLENCE IN IRAQ*

Georgia Papadogeorgou† Kosuke Imai‡ Jason Lyall§ Fan Li¶

June 9, 2022

Abstract

Many causal processes have spatial and temporal dimensions. Yet the classic causal inference framework is not directly applicable when the treatment and outcome variables are generated by spatio-temporal point processes. We extend the potential outcomes framework to these settings by formulating the treatment point process as a stochastic intervention. Our causal estimands include the expected number of outcome events in a specified area under a particular stochastic treatment assignment strategy. Our methodology allows for arbitrary patterns of spatial spillover and temporal carryover effects. Using martingale theory, we show that the proposed estimator is consistent and asymptotically normal as the number of time periods increases. We propose a sensitivity analysis for the possible existence of unmeasured confounders, and extend it to the Hájek estimator. Simulation studies are conducted to examine the estimators' finite sample performance. Finally, we illustrate the proposed methods by estimating the effects of American airstrikes on insurgent violence in Iraq from February 2007 to July 2008. Our analysis suggests that increasing the average number of daily airstrikes for up to one month may result in more insurgent attacks. We also find some evidence that airstrikes can displace attacks from Baghdad to new locations up to 400 kilometers away.

Keywords: carryover effects, inverse probability of treatment weighting, point process, sensitivity analysis, spillover effects, stochastic intervention, unstructured interference

*This material is based upon work partially supported by the National Science Foundation under Grant No. 2124124, 2124463, and 2124323. Lyall gratefully acknowledges financial support from the Air Force Office of Scientific Research (Grant #FA9550-14-1-0072). The findings and conclusions reached here do not reflect the official views or policy of the United States Government or Air Force. In addition, Imai thanks the Sloan Foundation (# 2020–13946) for financial support. The authors would also like to thank Soubhik Barari, Iavor Bojinov, Naoki Egami, Connor Jerzak, Sayar Karmakar, and Neil Shephard for their constructive comments.

†Assistant Professor, Department of Statistics, University of Florida, Gainesville FL 32611. Email: gpapadogeorgou@ufl.edu, URL: https://gpapadogeorgou.netlify.com

‡Professor, Department of Government and Department of Statistics, Harvard University. 1737 Cambridge Street, Institute for Quantitative Social Science, Cambridge MA, 02138. Email: imai@Harvard.Edu, URL: https://imai.fas.harvard.edu

§James Wright Chair in Transnational Studies and Associate Professor, Department of Government, Dartmouth College, Hanover, NH 03755. Email: jason.lyall@dartmouth.edu, URL: www.jasonlyall.com

¶Professor, Department of Statistical Science, Duke University, Durham, NC 27708. Email: fl35@duke.edu, URL: http://www2.stat.duke.edu/~fl35

1 Introduction

Many causal processes involve both spatial and temporal dimensions. Examples include the environmental impact of newly constructed factories, the economic and social effects of refugee flows, and the various consequences of disease outbreaks. These applications also illustrate key methodological challenges. First, when the treatment and outcome variables are generated by spatio-temporal processes, there exists an infinite number of possible treatment and event locations at each point in time. In addition, spatial spillover and temporal carryover effects are likely to be complex and may not be well understood.

Unfortunately, the classical causal inference framework that dates back to Neyman (1923) and Fisher (1935) is not directly applicable to such settings. Indeed, standard causal inference approaches assume that the number of units that can receive the treatment is finite (e.g., Rubin, 1974; Robins, 1997). Although a small number of studies develop a continuous time causal inference framework, they do not incorporate a spatial dimension (e.g., Gill and Robins, 2001; Zhang *et al.*, 2011). In addition, causal inference methods have been used for analyzing functional magnetic resonance imaging (fMRI) data, which have both spatial and temporal dimensions. For example, Luo *et al.* (2012) apply randomization-based inference, while Sobel and Lindquist (2014) employ structural modelling. We instead focus on data generated by different underlying processes, leading to new estimands and estimation strategies.

Specifically, we consider settings in which the treatment and outcome events are assumed to be generated by spatio-temporal point processes (Section 3). The proposed method is based on a single time series of spatial patterns of treatment and outcome variables, and builds upon three strands of the causal inference literature: interference, stochastic interventions, and time series.

First, we address the possibility that treatments might affect outcomes at a future time period and at different locations in arbitrary ways. Although some researchers have considered unstructured interference, they assume non-spatial and cross-sectional settings (see e.g., Basse and Airoldi, 2018; Sävje *et al.*, 2019, and references therein). In addition, Aronow *et al.* (2019) study spatial randomized experiments in a cross-sectional setting, and under the assumption that the number of potential intervention locations is finite and their spatial coordinates are known and fixed. By contrast, our proposed spatio-temporal causal inference framework allows for *temporally and spatially unstructured interference* over an infinite number of locations.

Second, instead of separately estimating the causal effects of treatment received at each location, we consider the impacts of different *stochastic treatment assignment strategies*, defined formally as the intervention distributions over treatment point patterns. Stochastic interventions have been used to estimate effects of realistic treatment assignment strategies (Díaz Muñoz and van der Laan, 2012; Young *et al.*, 2014; Papadogeorgou *et al.*, 2019) and to address challenging causal inference problems including violation of the positivity assumption (Kennedy, 2019), interference (Hudgens and Halloran, 2008; Imai *et al.*, 2021), mediation analysis (Lok, 2016; Díaz and Hejazi, 2019), and multiple treatments (Imai and Jiang, 2019). We show that this approach is also useful for causal inference with spatio-temporal treatments and outcomes.

Finally, our methodology allows for arbitrary patterns of spatial and temporal interference. As such, our estimation method does not require the separation of units into minimally interacting sets (e.g., Tchetgen Tchetgen *et al.*, 2017). Nor does it rely on an outcome modelling approach that entails specifying a functional form of spillover effects based on, for example, geographic distance. Instead, we view our data as a single time series of maps, which record the locations of treatment and outcome realizations as well as the geographic coordinates of other relevant events. Our estimation builds on the time-series causal inference approach pioneered by Bojinov and Shephard (2019).

We propose a spatially-smoothed inverse probability weighting estimator that is consistent and asymptotically normal under a set of reasonable assumptions, regardless of whether the propensity scores are known, or estimated from a correctly specified model (Section 4). To do so, we establish a new central limit theorem for martingales that can be widely used for causal inference in observational, time series settings. We also show that the proposed estimator based on the estimated propensity score has a lower asymptotic variance than when the true propensity score is known. This generalizes the existing theoretical result under the independently and identically distributed setting (Hirano *et al.*, 2003) to the spatially and temporally dependent setting. Finally, to assess the potential impact of unobserved confounding, we develop a sensitivity analysis method by generalizing the sensitivity analysis of Rosenbaum (2002) to our spatio-temporal context and to the Hájek estimator with standardized weights (Section 5). We conduct simulation studies to assess the finite sample performance of the proposed estimators (Section 6).

Our motivating illustration is the evaluation of the effects of American airstrikes on insurgent violence in Iraq from February 2007 to July 2008 (Section 2). We consider all airstrikes during each day anywhere in Iraq as a *treatment pattern*. Instead of focusing on the causal effects of each airstrike, we estimate the effects of different *airstrike strategies*, defined formally as the distributions of airstrikes throughout Iraq (Section 7). The proposed methodology enables us to capture spatio-temporal variations in treatment effects, shedding new light on how airstrikes affect the location, distribution, and intensity of insurgent violence.

Specifically, under a set of assumptions, our analysis suggests that a higher number of airstrikes, without modifying their spatial distribution, may increase the number of insurgent attacks, especially near Baghdad, Mosul, and the roads between them. We also find that changing the focal point of airstrikes to Baghdad without modifying the overall frequency can shift insurgent attacks from Baghdad to Mosul and its environs. Under our assumptions, these findings suggest that airstrikes can increase insurgent attacks *and* disperse them over considerable distances. Furthermore, our analysis shows that increasing the number of airstrikes may initially reduce attacks but ultimately increase them over the long run. Our sensitivity analysis indicates, however, that these findings are somewhat sensitive to the potential existence of unmeasured confounders. Thus, further analyses are necessary in order for us to reach more definitive conclusions about the impacts of airstrikes.

The proposed methodology has a wide range of applications beyond the specific example analyzed in this paper. For example, the causal effects of pandemics and crime on a host of economic and social outcomes could be evaluated using our methodology. With the advent of massive and granular data sets, we expect the need to conduct causal analysis of spatio-temporal data will only continue to grow.

2 Motivating Application: Airstrikes and Insurgent Activities in Iraq

Airstrikes have emerged as a principal tool for fighting against insurgent and terrorist organizations in civil wars around the globe. In the past decade alone, the United States has conducted sustained air campaigns in at least six different countries, including Afghanistan, Iraq, and Syria. Although it has been shown that civilians have all-too-often borne the brunt of these airstrikes (Lyall, 2019b), we have few rigorous studies that evaluate the impact of airstrikes on subsequent insurgent violence. Even these studies have largely reached opposite conclusions, with some claiming that airpower reduces insurgent attacks while others arguing they spark escalatory spirals of increased violence (e.g., Lyall, 2019a; Mir and Moore, 2019a; Dell and Querubin, 2018; Kocher *et al.*, 2011).

Moreover, all existing studies have two interrelated methodological shortcomings: they carve continuous geographic space into discrete, often arbitrary, units, and they make simplifying assumptions about patterns of spatial and temporal interference. Mir and Moore (2019b), for example, argue that drone strikes in Pakistan have reduced terrorist violence. But they use a coarse estimation strategy that bins average effects of drone strikes into broad half-year increments over entire districts that cannot capture local spatial and temporal dynamics. Similarly, Rigterink (2021) draws on 443 drone strikes to estimate airstrike effects on 13 terrorist groups in Pakistan, concluding that they have mixed effects. Yet her group-month estimation strategy cannot detect spillover effects nor accurately capture the timing of insurgent responses. In short, we need a flexible methodological approach that avoids the pitfalls of binning treatment and outcome measures into too-aggregate, possibly misleading, temporal and spatial units.

We enter this debate by examining the American air campaign in Iraq. We use declassified US Air Force data on airstrikes and shows of force (simulated airstrikes where no weapons are released) for the February 2007 to July 2008 period. The period in question coincides with the "surge" of American forces and airpower designed to destroy multiple Sunni and Shia insurgent organizations in a bid to turn the war's tide.

Aircraft were assigned to bomb targets via two channels. First, airstrikes were authorized in response to American forces coming under insurgent attack. These close air support missions represented the vast majority of airstrikes in 2007–08. Second, a small percentage (about 5%) of airstrikes were pre-planned against high-value targets, typically insurgent commanders, whose presence had been detected from intercepted communications or human intelligence. In each case, airstrikes were driven by insurgent attacks that were either ongoing or had occurred in the recent past in a given location. As a result, the models used later in this paper adjust for prior patterns of insurgent violence in a given location for several short-term windows.

We also account for prior air operations, including shows of force, by American and allied aircraft. Insurgent violence in Iraq is also driven by settlement patterns and transportation networks. Our models therefore include population size and location of Iraqi villages and cities as well as proximity to road networks, where the majority of insurgent attacks were conducted against American convoys. Finally, prior reconstruction spending might also drive the location of airstrikes. Aid is often provided in tandem with airstrikes to drive out insurgents, while these same insurgents often attack aid locations to derail American hearts-and-minds strategies. Taken together, these four factors—recent insurgent attacks, the presence of American forces, settlement patterns, and prior aid spending—drove decisions about the location and severity of airstrikes. We emphasize that we may not observe all factors used for decisions on airstrikes. We will address this limitation by developing and applying a sensitivity analysis.

Figure 1 summarizes the spatial and temporal distributions of airstrikes (treatment variable) and insurgent violence (outcome variable). Figure 1a presents the temporal distribution of airstrikes recorded by the US Air Force each month. There were a total of 2,246 airstrikes during this period. Figure 1b plots the spatial density of these airstrikes across Iraq, with spatial clustering observed around Baghdad