Causal Inference with Spatio-temporal Data: Estimating the Effects of Airstrikes on Insurgent Violence in Iraq

Kosuke Imai

Harvard University

Talk at European University Institute
May 20, 2020

Joint work with Georgia Papadogeorgou (Duke),
Jason Lyall (Dartmouth), and Fan Li (Duke)

Motivation

- Increasing availability of unstructured data in social sciences
    - don't come in a nice matrix form $\leftrightsquigarrow$ survey, official statistics
    - text, images, audio, video, etc.
- How should we draw causal inference from these new types of data?
- Causal inference with spatio-temporal data
    - a time series of maps as data
    - treatment and outcome event locations in a continuous space
    - applications: crime incidents, disease outbreaks, etc.
- Methodological challenges
    1. spillover effects over space
    2. carryover effects over time
    3. infinitely many possible treatment and outcome locations
- Current practice
    1. arbitrary discretization of space
    2. assumptions about spillover and carryover effects

Impacts of Airstrikes on Insurgent Violence in Iraq

- Airstrikes as a principal tool for combating insurgency in civil wars
- Debate: whether or not airstrikes reduce subsequent insurgent attacks
  (e.g., Kocher *et al.* 2011; Dell and Querubin 2018; Lyall 2019; Mir and Moore 2019)
- Methodological limitations:
  - discretize continuous space into aggregate geographical units
  - simplifying assumptions about spillover and carryover effects

- American air campaign in Iraq:
  - declassified USAF data from Jan. 2007 to July 2008 ("surge" period)
  - daily data with precise geolocation for airstrikes and insurgent attacks
- Drivers of airstrikes:
  - prior patterns of insurgent attacks and airstrikes
  - presence of American forces
  - settlement patterns and road networks
  - economic aid
  - intelligence about high-value targets (small fraction)

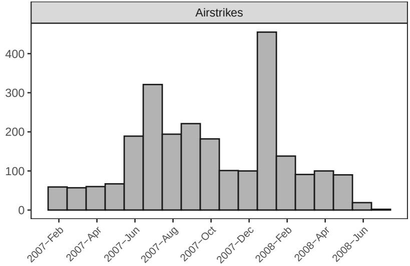

(a) airstrikes over time

(b) Airstrikes over space

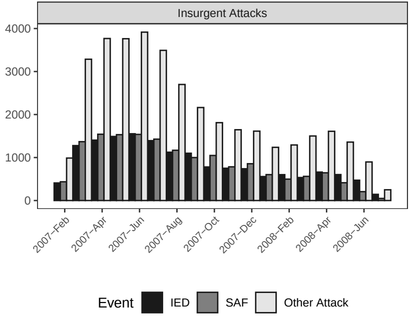

(c) Insurgent violence over time

(d) Insurgent violence over space

Contributions

- Causal inference with point process treatment and outcome
  - impossible to estimate causal effects of each treatment event
    - unrestricted spillover and carryover effects
    - probability of each treatment realization is zero $\rightsquigarrow$ lack of overlap
  - stochastic intervention based on the distribution of treatments
  - distribution of airstrikes as a military strategy
- Causal estimands under stochastic intervention
  - expected number of outcome events within a region of interest
  - various stochastic interventions
    1. change the dosage while keeping the distribution identical
    2. change the distribution while keeping the overall dosage constant
    3. intervention over multiple time periods
- The proposed IPW (inverse probability of treatment) estimator
  - overlap and unconfoundedness assumptions
  - consistency and asymptotic normality
- Simulation studies and empirical application

The Setup

- T time periods: $t = 1, 2, \ldots, T$
- Treatment variable
  - Ω: set of all possibly infinite locations that can receive the treatment
  - $W_t(s) \in \{0, 1\}$: binary treatment indicator for location s at time t
  - $W_t = \{W_t(s) : s \in \Omega\} \in \mathcal{W}$: treatment location map at time t
  - $S_{W_t} = \{s \in \Omega : W_t(s) = 1\}$: set of treatment-active locations at time t with $|S_{w_t}| < \infty$
  - $\overline{\boldsymbol{W}}_t = (W_1, W_2, \ldots, W_t)$: observed treatment history up to time t
- Outcome variable
  - $Y_t(s)$, Y_t, and $\overline{\boldsymbol{Y}}_t$ can be similarly defined
  - Potential outcome: $Y_t(\overline{\boldsymbol{w}}_t)$ where $w_t \in \mathcal{W}$ is a realized treatment and $\overline{\boldsymbol{w}}_t = (w_1, w_2, \ldots, w_t) \in \mathcal{W}^t$ is a treatment history realization at time t
  - Observed outcome: $Y_t = Y_t(\overline{\boldsymbol{W}}_t)$
  - $S_{Y_t(\overline{\boldsymbol{w}}_t)}$: set of outcome-active locations under treatment history $\overline{\boldsymbol{w}}_t$
  - History of all potential outcomes up to time t: $\overline{\mathcal{Y}}_t = \{Y_{t'}(\overline{\boldsymbol{w}}_{t'}) : \overline{\boldsymbol{w}}_{t'} \in \mathcal{W}^{t'}, t' \leq t\}$
- Time-varying confounders: X_t, $\overline{\mathbf{X}}_t$, $X_t(\overline{\boldsymbol{w}}_{t-1})$, and $\overline{\mathcal{X}}_t$

Stochastic Intervention

- Stochastic intervention: any distribution of treatment can be used
- We consider Poisson point process F_h
  - homogeneous Poisson point process with intensity h:
    For any disjoint region $B_1, B_2, \ldots, B_n \subset \Omega$, the number of events in each region B_i

$$N_{B_i}(W) \overset{\text{indep.}}{\sim} \text{Poisson}(h|B_i|)$$

  - non-homogeneous Poisson point process with intensity function $h(\omega)$:

$$N_{B_i}(W) \overset{\text{indep.}}{\sim} \text{Poisson}\left(\int_{B_i} h(\omega)d\omega\right)$$

- Example:

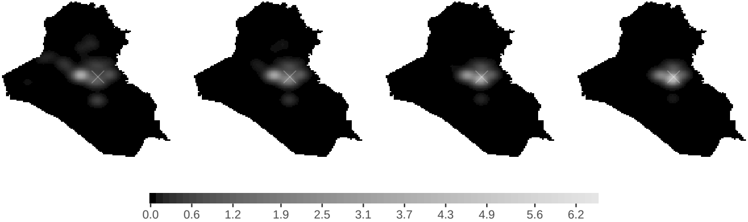

Causal Estimands

- Expected number of outcome-active locations in region B at time t under stochastic intervention F_h conducted at time t

$$\overline{N}_{Bt}(F_h) \ = \ \int_{\mathcal{W}} N_B(Y_t(\overline{\boldsymbol{W}}_{t-1}, w_t))dF_h(w_t)$$

- Further average this quantity over time:

$$\overline{N}_B(F_h) \ = \ \frac{1}{T}\sum_{t=1}^{T} \overline{N}_{Bt}(F_h)$$

- We can compare the different interventions:

$$\tau_B(F_{h'}, F_h) \ = \ \overline{N}_B(F_{h'}) - \overline{N}_B(F_h)$$

Stochastic Intervention over Multiple Time Periods

- Consider a non-dynamic stochastic intervention over M time periods

$$F_{\boldsymbol{h}} \;=\; F_{h_1} \times \cdots \times F_{h_M} \quad \text{where } \boldsymbol{h} = (h_1, h_2, \ldots, h_M)$$

- Expected number of outcome-active locations in region B at time t under stochastic intervention $F_{\boldsymbol{h}}$ conducted from time $t - M + 1$ to t

$$\overline{N}_{Bt}(F_{\boldsymbol{h}}) \;=\; \int_{\mathcal{W}} \cdots \int_{\mathcal{W}} N_B(Y_t(\overline{\boldsymbol{W}}_{t-M}, w_{t-M+1}, \ldots, w_t)) \\ dF_{h_M}(w_{t-M+1}) \cdots dF_{h_1}(w_t)$$

- Average this quantity over time:

$$\overline{N}_B(F_{\boldsymbol{h}}) \;=\; \frac{1}{T - M + 1} \sum_{t=M}^{T} \overline{N}_{Bt}(F_{\boldsymbol{h}})$$

- Comparison of different interventions:

$$\tau_B(F_{\boldsymbol{h}'}, F_{\boldsymbol{h}}) \;=\; \overline{N}_B(F_{\boldsymbol{h}'}) - \overline{N}_B(F_{\boldsymbol{h}})$$

e.g., lagged effects with $h_M \neq h'_M$ and $\boldsymbol{h}_{-M} = \boldsymbol{h}'_{-M}$

Assumptions

1. Unconfoundedness: treatment is independent of all potential (past and future) paths for the outcome and time-varying confounders conditional on the observed history

$$f(W_t \mid \overline{\boldsymbol{W}}_{t-1}, \overline{\boldsymbol{Y}}_{t-1}, \overline{\boldsymbol{X}}_t, \{\overline{\mathcal{Y}}_T, \overline{\mathcal{X}}_T\}) \;=\; f(W_t \mid \overline{\boldsymbol{W}}_{t-1}, \overline{\boldsymbol{Y}}_{t-1}, \overline{\boldsymbol{X}}_t)$$

$\rightsquigarrow$ the generalization of the non-anticipating assumption for time-series experiments (Bojinov and Shephard, 2019)

2. Overlap: there exists a constant $\delta_W > 0$ such that

$$\underbrace{f(W_t = w \mid \overline{\boldsymbol{W}}_{t-1}, \overline{\boldsymbol{Y}}_{t-1}, \overline{\boldsymbol{X}}_t)}_{propensity\ \ score} \;>\; \delta_W \cdot \underbrace{f_h(w)}_{density\ \ of\ \ F_h} \qquad \text{for all } w \in \mathcal{W}$$

$\rightsquigarrow$ the ratio $f_h(w)/f(W_t = w \mid \overline{\boldsymbol{W}}_{t-1}, \overline{\boldsymbol{Y}}_{t-1}, \overline{\boldsymbol{X}}_t)$ is bounded

The Proposed Estimator

- Inverse probability of treatment weighting (IPW)
- Kernel smoothing of spatial point patterns
- Estimated outcome surface at $\omega \in \Omega$

$$\widehat{Y}_t(F_h;\omega) \; = \; \underbrace{\frac{\overbrace{f_h(W_t)}^{counterfactual \;\; distribution}}{\hat{f}(W_t \mid \overline{\boldsymbol{W}}_{t-1}, \overline{\boldsymbol{Y}}_{t-1}, \overline{\boldsymbol{X}}_t)}}_{actual \;\; distribution} \; \underbrace{\sum_{s \in S_{Y_t}} K_b(\|\omega - s\|)}_{spatially \;\; smoothed \;\; outcome}$$

where K_b is the scaled Kernel function with bandwidth parameter b

- Estimated number of outcome-active locations in region B

$$\widehat{\widehat{N}}_{Bt}(F_h) \; = \; \int_B \widehat{Y}_t(F_h;\omega) d\omega$$

- Averaging over time

$$\widehat{\widehat{N}}_B(F_h) \; = \; \frac{1}{T} \sum_{t=1}^{T} \widehat{\widehat{N}}_{Bt}(F_h)$$

Estimation for Intervention over Multiple Time Periods M

- Estimated outcome surface at $\omega \in \Omega$

$$\widehat{Y}_t(F_{\boldsymbol{h}};\omega) \;=\; \underbrace{\prod_{j=t-M+1}^{t} \frac{f_{h_{t-j+1}}(W_j)}{\hat{f}(W_j \mid \overline{\boldsymbol{W}}_{j-1}, \overline{\boldsymbol{Y}}_{j-1}, \overline{\boldsymbol{X}}_j)}}_{product\ \ of\ \ M\ \ ratios} \sum_{s \in S_{Y_t}} K_b(\|\omega - s\|)$$

  where K_b is the scaled Kernel function with bandwidth parameter b
- Estimated number of outcome-active locations in region B

$$\widehat{\widehat{N}}_{Bt}(F_{\boldsymbol{h}}) \;=\; \int_B \widehat{Y}_t(F_{\boldsymbol{h}};\omega) d\omega$$

- Averaging over time

$$\widehat{\widehat{N}}_B(F_{\boldsymbol{h}}) \;=\; \frac{1}{T-M+1} \sum_{t=M}^{T} \widehat{\widehat{N}}_{Bt}(F_{\boldsymbol{h}})$$

Asymptotic Normality and Variance Estimation

- Suppose that there exists v such that

$$\frac{1}{T-M+1}\sum_{t=M}^{T} v_t \xrightarrow{p} v \quad \text{as } T \longrightarrow \infty$$

where $v_t = \mathbb{V}\left(\widehat{\overline{N}}_{Bt}(F_{\boldsymbol{h}}) \mid \overline{\boldsymbol{W}}_{t-M}, \overline{\mathcal{Y}}_T, \overline{\mathcal{X}}_T\right)$.

- Then, under some regularity conditions, we have,

$$\sqrt{T}\left(\widehat{\overline{N}}_B(F_{\boldsymbol{h}}) - \overline{N}_B(F_{\boldsymbol{h}})\right) \xrightarrow{d} \mathcal{N}(0, v)$$

where the proof is based on the martingale theory

- Time-specific variance v_t cannot be estimated
- We use the upper bound $v_t \leq \mathbb{E}\left(\widehat{\overline{N}}_{Bt}(F_{\boldsymbol{h}})^2 \mid \overline{\boldsymbol{W}}_{t-M}, \overline{\mathcal{Y}}_T, \overline{\mathcal{X}}_T\right)$
- Estimated propensity score $\rightsquigarrow$ smaller variance

Empirical Analysis: Setup

- Estimate the baseline treatment distribution $\phi_0(\omega)$ based on the airstrikes data from January to September, 2006

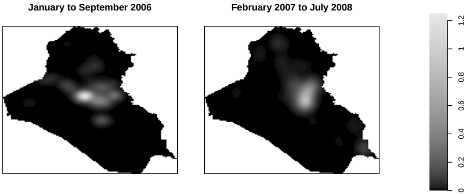

1. How does increasing airstrikes affect insurgent violence?
   $\rightsquigarrow$ vary $c > 0$ for $h(\omega) = c \cdot \phi_0(\omega)$
2. How long does it take for these effects to be realized?
   $\rightsquigarrow$ vary c for $h_M(\omega) = c \cdot \phi_0(\omega)$ and $h_1(\omega) = \cdots = h_{M-1}(\omega) = \phi_0(\omega)$
3. How does the shift in the prioritization of certain locations for airstrikes change the spatial pattern of insurgent attacks?
   $\rightsquigarrow$ vary $\alpha > 0$ for $h_\alpha(\omega) = c_\alpha \cdot \phi_0(\omega) d_\alpha(\omega)$ with $\int_\Omega h_\alpha(\omega) d\omega = c$
   - power density $d_\alpha(\omega) \propto d(\omega)^\alpha$
   - $d(\omega) =$ the normal density centered at s_f with precision α

Propensity Score Model Specification

- Non-homogeneous Poisson point process
  - prior airstrikes over the last day, week, and month

$$\overline{W}^*_{t-1}(\omega) = \sum_{j=1}^{7} \sum_{s \in S_{W_{t-j}}} \exp\{-\|s-\omega\|\}$$

  - prior insurgent attacks over the last day, week, and month
  - prior show-of-force over the last day, week, and month
  - amount of US aid in each district over the past month
  - distances from major cities, road networks, rivers, and settlements
  - log population of governorate (measured in 2003), temporal splines

Increasing the Expected Number of Airstrikes from 1 to 6 per Day Leads to More Insurgent Attacks

(a) Improvised Explosive Device (IED)

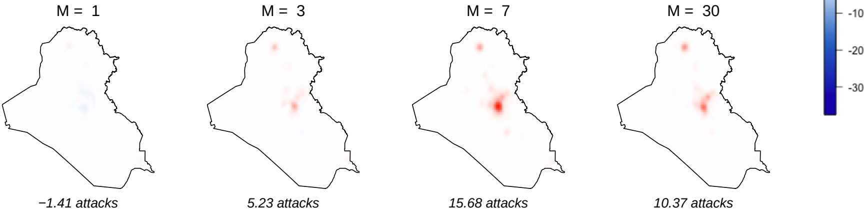

(b) Small Armed Fire (SAF)

Effect of Increasing the Airstrikes **for M Days** on the Number of Insurgent Attacks within Baghdad

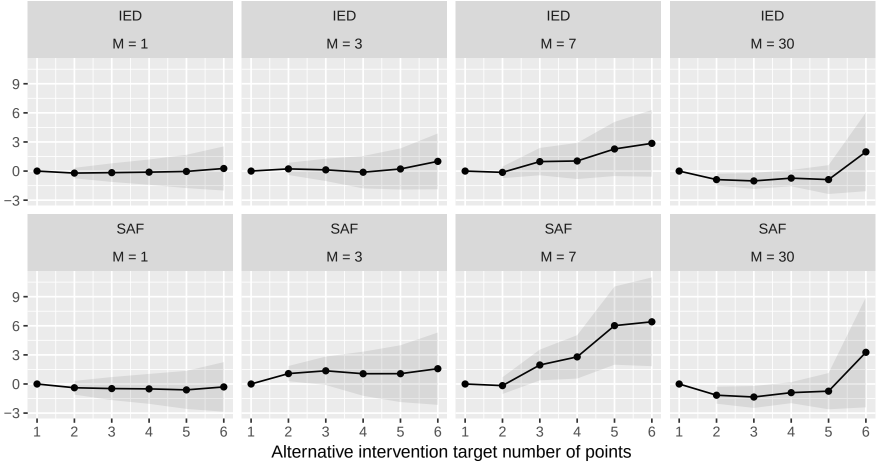

Effect of Increasing the Airstrikes M **Days Ago** on the Number of Insurgent Attacks within Baghdad

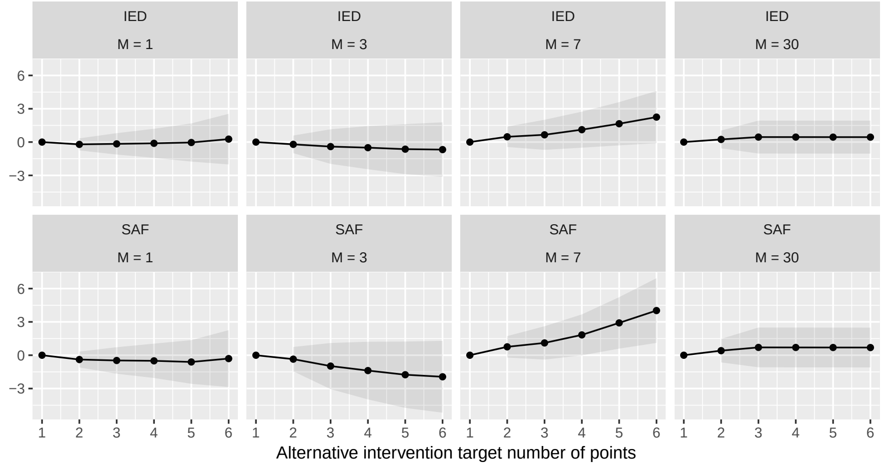

Increasing the Priority of Baghdad as a Focal Point of Airstrikes Shifts Attacks to Mosul when *M* is Large

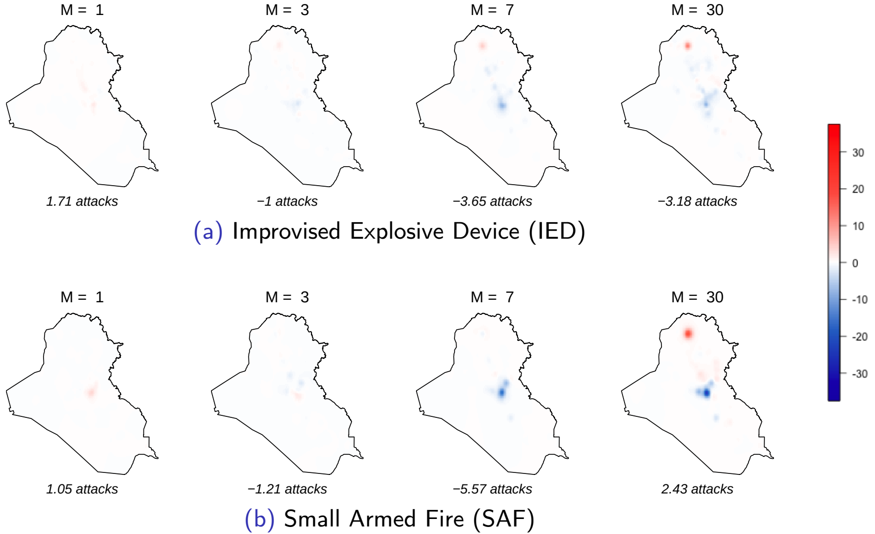

Concluding Remarks

- A new approach to causal inference with spatio-temporal data
  - directly model point patterns without arbitrary aggregation
  - allow for unstructured spillover and carryover effects
- Key idea: stochastic intervention
  - consider treatment distributions rather than fixed treatment values
  - can handle infinitely many possible treatment locations
  - combine this with spatial smoothing for outcome point process
- Effects of airstrikes on insurgent attacks in Iraq
  - airstrike strategies as stochastic interventions
  - flexible estimation of spillover and carryover effects
- Future research:
  - causal inference with unstructured data such as texts
  - civilian casualty as mediator; comparison with hearts and minds
- Paper at https://imai.fas.harvard.edu/research/spatiotempo.html