

Optional ERIC Coversheet — Only for Use with U.S. Department of Education Grantee Submissions

This coversheet should be completed by grantees and added to the PDF of your submission if the information required in this form **is not included on the PDF to be submitted**.

INSTRUCTIONS

- Before beginning submission process, download this PDF coversheet if you will need to provide information not on the PDF.
- Fill in all fields—information in this form **must match** the information on the submitted PDF and add missing information.
- Attach completed coversheet to the PDF you will upload to ERIC [use Adobe Acrobat or other program to combine PDF files]—do not upload the coversheet as a separate document.
- Begin completing submission form at <https://eric.ed.gov/submit/> and upload the full-text PDF with attached coversheet when indicated. Your full-text PDF will display in ERIC after the 12-month embargo period.

GRANTEE SUBMISSION REQUIRED FIELDS

Title of article, paper, or other content

All author name(s) and affiliations on PDF. If more than 6 names, ERIC will complete the list from the submitted PDF.

Last Name, First Name	Academic/Organizational Affiliation	ORCID ID

Publication/Completion Date—(if *In Press*, enter year accepted or completed)

Check type of content being submitted and complete one of the following in the box below:

- If article: Name of journal, volume, and issue number if available
- If paper: Name of conference, date of conference, and place of conference
- If book chapter: Title of book, page range, publisher name and location
- If book: Publisher name and location
- If dissertation: Name of institution, type of degree, and department granting degree

DOI or URL to published work (if available)

Acknowledgement of Funding— Grantees should check with their grant officer for the preferred wording to acknowledge funding. If the grant officer does not have a preference, grantees can use this suggested wording (adjust wording if multiple grants are to be acknowledged). Fill in Department of Education funding office, grant number, and name of grant recipient institution or organization.

“This work was supported by U.S. Department of Education [Office name]
through [Grant number] to Institution] . The opinions expressed are
those of the authors and do not represent views of the [Office name]
or the U.S. Department of Education.

ESTIMATING THE EFFECTS OF A CALIFORNIA GUN CONTROL PROGRAM WITH MULTITASK GAUSSIAN PROCESSES

BY ELI BEN-MICHAEL¹, DAVID ARBOUR², AVI FELLER^{3,*}, ALEX FRANKS⁴, AND STEVEN RAPHAEL³

¹*Carnegie Mellon University*

²*Adobe Research*

³*UC Berkeley* *afeller@berkeley.edu

⁴*UC Santa Barbara*

Gun violence is a critical public safety concern in the United States. In 2006 California implemented a unique firearm monitoring program, the Armed and Prohibited Persons System (APPS), to address gun violence in the state. The APPS program first identifies those firearm owners who become prohibited from owning one due to federal or state law, then confiscates their firearms. Our goal is to assess the effect of APPS on California murder rates using annual, state-level crime data across the US for the years before and after the introduction of the program. To do so, we adapt a non-parametric Bayesian approach, multitask Gaussian Processes (MTGPs), to the panel data setting. MTGPs allow for flexible and parsimonious panel data models that nest many existing approaches and allow for direct control over both dependence across time and dependence across units, as well as natural uncertainty quantification. We extend this approach to incorporate non-Normal outcomes, auxiliary covariates, and multiple outcome series, which are all important in our application. We also show that this approach has attractive Frequentist properties, including a representation as a weighting estimator with separate weights over units and time periods. Applying this approach, we find that the increased monitoring and enforcement from the APPS program substantially decreased homicides in California. We also find that the effect on murder is driven entirely by declines in gun-related murder with no measurable effect on non-gun murder. Estimated cost per murder avoided are substantially lower than conventional estimates of the value of a statistical life, suggesting a very high benefit-cost ratio for this enforcement effort.

1. Introduction. Every year, thousands of people in the United States are murdered with a firearm. In 2020 alone, there were over 45,000 firearm-related deaths in the U.S., including deaths by suicide and murder; firearm-related injuries are now the leading cause of death among U.S. children and adolescents ([Goldstick, Cunningham and Carter, 2022](#)). This mortality rate remains multiple times that of other wealthy nations, largely due to differences in the lethality of violence in the United States (see [Hemenway, 2017](#); [Zimring, 1972](#)).

Addressing gun violence is thus one of the most pressing public safety issues in the United States. Policy proposals intended to curb gun violence vary considerably. At one end of the spectrum are proposals to increase gun availability to deter potential assailants through such policies as shall-issue concealed carry laws (see [Donohue, Aneja and Weber, 2019](#)) or the arming of school teachers ([Green and Fernandez, 2018](#)). At the other end are policy proposals to restrict access to firearms through universal background checks, temporarily removing firearms from individuals deemed a danger to themselves or others ([Wintemute et al., 2019](#)), and restricting access to assault weapons with high capacity magazines ([Donohue and Boulouta, 2019](#)).

Keywords and phrases: Gun violence, Gaussian Processes, Panel Data, Causal Inference.

Federal and state laws already prohibit certain individuals from owning firearms. Under federal law, individuals with prior felony convictions as well as some individuals with a history of mental illness are prohibited from purchasing and owning firearms. Many states go further, restricting access to firearms for individuals convicted of violent misdemeanors, individuals that are subjected to restraining orders, and juvenile offenders, among other categories ([Giffords Law Center to Prevent Gun Violence, 2019](#)). While law enforcement officials confiscate firearms of prohibited persons discovered to be in possession, until recently there was no effort to proactively and comprehensively recover firearms from prohibited individuals.

California was the first and remains the only state to undertake such an effort. In 2006 California implemented the Armed and Prohibited Persons System (APPS), a monitoring program that identifies known firearm owners who become prohibited from owning a firearm and then subsequently attempts to retrieve all prohibited weapons. Through this program, state law enforcement has made hundreds of thousands of contacts with prohibited individuals and has removed tens of thousands of firearms from their possession, often permanently.¹

In this paper, we assess whether the implementation of APPS affected murder rates in California, using annual, state-level data for the 50 states. To do so, we adapt a non-parametric Bayesian approach, *multitask Gaussian Processes* (MTGPs), to estimate the causal effect of APPS. Originally developed for the setting with multiple outcomes rather than multiple units ([Bonilla, Chai and Williams, 2008](#); [Álvarez, 2017](#)), Multitask GPs separately parameterize dependence across time and dependence across units. This allows for a flexible and parsimonious panel data model that nests many existing approaches: Our primary model is a natural generalization of low-rank factor models commonly used in panel data settings (e.g., [Xu, 2017](#); [Athey et al., 2021](#)) with an additional prior that encourages smoothness in the underlying factors. We also take advantage of the large literature on the statistical and computational properties of Gaussian Processes ([Williams and Rasmussen, 2006](#); [Gelman et al., 2013](#)) to extend this model. Immediate extensions include: allowing for a count (Poisson) observational model; allowing for multiple outcome series; incorporating a mean model; and incorporating auxiliary covariates.

Applying these ideas to the impact of APPS yields large, negative effect estimates on gun-related homicides in California. Specifically, we estimate over 100 murders prevented per year, a decline of at least 5 percent relative to California’s murder rate immediately prior to the program. In addition, we find that the effect on murder is driven entirely by declines in gun-related murder with no measurable effect on non-gun murder. Back-of-the-envelope calculations find that program expenditures are likely under \$100,000 per murder prevented, much lower than conventional estimates of the value of a statistical life, suggesting a very high benefit-cost ratio for this enforcement effort.

Methodologically, we believe that the MTGP framework is a promising approach for estimating causal effects with panel data more broadly. First, we can naturally incorporate many panel data diagnostics and other model checks in terms of standard Bayesian workflows, such as posterior predictive checks. This is especially important in settings such as ours, with relatively few units and (pre-treatment) time periods. The Bayesian approach also yields coherent

¹While California is not the only state that screens existing gun owners for prohibiting conditions, it is the only state to screen daily for triggering events, to carefully monitor whether the person has been disarmed, and to have a law enforcement unit in charge of following up, disarming such individuals, and storing the firearms. That being said, 19 states and D.C. currently have Extreme Risk Protection Orders ([Giffords Law Center to Prevent Gun Violence, 2022](#)) and several local jurisdictions have targeted efforts to remove firearms from those who are a danger to themselves or others as well as to prevent straw purchases of firearms for those who cannot legally purchase. For an example of the former, see [Barron \(2022\)](#). For an example of the latter, see the description of the Chicago Anti-Gun Enforcement Program, in operation during the 1990s ([Sheppard, 1999](#)).

uncertainty quantification, appropriately propagating uncertainty from the different sources of information in the model. In particular, the specific model we consider automatically produces uncertainty intervals that grow over time post-treatment. Surprisingly, this last point is a departure from many existing approaches, which sometimes have constant uncertainty intervals post-treatment. We also report many flexible summaries of the posterior distribution, allowing us to obtain posterior distributions for estimands like the benefit-cost ratio, a key quantity in policy analysis.

Finally, we adapt results from [Kanagawa et al. \(2018\)](#) showing that the Gaussian Process can be represented as a weighting estimator. We then use this representation to show that the overall weights decompose into separate unit and time weights. As a result, the proposed MTGP approach has the form of a doubly-weighted average over units and time periods, similar to recent proposals in the panel data literature ([Arkhangelsky et al., 2019](#); [Ben-Michael, Feller and Rothstein, 2021](#)).

As we discuss in [Section 3.3](#), our proposed approach combines the growing literature on Gaussian Processes for causal inference (see [Oganisian and Roy, 2020](#)) with the robust literature on estimating causal effects for a single treated unit with panel data (see [Samartsidis et al., 2019](#)). Of particular relevance are [Modi and Seljak \(2019\)](#) and [Carlson \(2020\)](#), who also consider GPs in this context, albeit with different structures. More broadly, our proposed framework is a natural (Bayesian) generalization of recent proposals for low-rank factor models with panel data (e.g., [Xu, 2017](#); [Athey et al., 2021](#)); see also recent extensions from [Pang, Liu and Xu \(2020\)](#).

Our paper proceeds as follows. In [Section 2](#) we give an overview of the APPS program and other institutional details and present an initial analysis of the effect of APPS. In [Section 3](#) we describe the underlying causal problem and review related work. In [Section 4](#) we review single- and multi-task Gaussian Processes and apply these ideas to our application setting. In [Section 5](#) we represent the MTGP approach as a weighting estimator. In [Section 6](#) we report model diagnostics and the estimated impact of APPS under different models. Finally, in [Section 7](#) we discuss open questions and possible extensions. The Appendix includes additional computational details, diagnostics, and results. All code and replication files are available at github.com/darbour/mtgp_panel.

2. Armed and Prohibited Persons System.

2.1. Description of the system. The Armed and Prohibited Persons System is a law enforcement program devoted to proactively removing firearms from known gun owners who become prohibited from owning a gun. The program has two components: an intricate data infrastructure and monitoring system that updates daily; and an enforcement operation, which includes both non-sworn analysts and a small force of sworn state officers, dedicated solely to enforcing state firearms restrictions for gun owners who become prohibited. While the state has maintained the database on gunowners and associated firearms since the mid 1990s, the active effort to remove firearms from prohibited individuals did not begin until December 2006 ([California Department of Justice, 2015](#)).

The core of the data system is the APPS database, a list of known gunowners that draws on Dealer Record of Sales (DROS) transactions, voluntary gun registrations with the state, and an Automated Firearm System (AFS) database that records all known firearms involved in sales through licensed dealers. The database then links each individual known to own a firearm to specific firearms identified by serial numbers. This system has recorded owners of handguns and their associated weapons since 1996. In 2014, the state expanded the system to include all gunowners and all firearms, including long guns. The known population of

gunowners in the APPS database has grown substantially over time, from under 900,000 in 2006 to over 2.5 million in 2019, or roughly 6 percent of the state's resident population.²

The APPS data are cross referenced daily against several databases that record events that trigger a firearm prohibition. These include the state's Automated Criminal History System (the state's criminal history repository), the Mental Health Reporting system (a database used by mental health providers to record involuntary commitments and other mental health events that may trigger a prohibition), the California Restraining and Protective Order system, and a statewide Wanted Persons System. When the cross-reference uncovers a triggering event for an individual in APPS, the case is analyzed by state criminal intelligence specialists to determine whether the person is actually prohibited and if the person is still armed. Local law enforcement across the state conducts gun relinquishment efforts for those who become prohibited; hence, local agencies may have already removed the person's firearms.³ In the event that the APPS database reveals that the person is still armed, the case is referred to the state Bureau of Firearms (BOF) for an enforcement action. Of these, 54 percent were prohibited due to a felony conviction, 11 percent were prohibited due to a misdemeanor conviction, 19 percent were prohibited due to an existing restraining order, 18 percent were prohibited due to a mental health prohibition, 24 percent were prohibited due to provisions of the Federal Brady Act, while the remaining 7 percent were prohibited due to active probation status ([California Department of Justice, 2019](#)).

BOF officers attempt to contact these individuals, usually at their last known address, and request to perform a consensual search of the premises. If refused, officers will return with a search warrant. In many instances, officers only recover some of the firearms associated with the individual in question, possibly due to a prior unrecorded sale, the individual turning over firearms to friends and family, the gun being lost or stolen, or local law enforcement already collecting the firearms at conviction and failing to update the person's record in APPS. In addition, officers frequently discover unregistered weapons along with the registered firearms.⁴ Once all guns associated with prohibited individuals are removed from possession, the individuals are removed from the APPS database as they are no longer armed. Gunowners are also removed from the APPS database upon death. Between 2007 and 2017, the BOF directly retrieved nearly 32,000 firearms from prohibited individuals. However, the total number of firearms relinquished statewide is likely much larger given the role of local law enforcement and the fact the BOF works cases that are somehow not handled locally.

2.2. Related research on access to firearms. A growing literature evaluates the impact of state-level changes in access to firearms on homicide rates; see [Webster and Wintemute \(2015\)](#) and [Cook and Donohue \(2017\)](#). Overall, recent studies find that more permissive firearm regulations, such as more permissive concealed carry laws and Right to Carry (RTC) laws, are associated with increases in gun-related deaths ([Donohue, Aneja and Weber, 2019](#); [Siegel et al., 2017](#)).⁵ Similarly, policy changes that restrict firearm access, such as permit-to-purchase laws and comprehensive background checks, are associated with decreases in

²The system likely covers only a portion of gunowners in the state since the underlying list of gunowners used by the system is generated largely by new firearms sales. For example, the list excludes people who legally purchase firearms in other states and fail to register them in California, as well as people who purchased handguns prior to 1996 and long guns prior to 2014 and have not registered these firearms.

³In 2016, this became the official responsibility of county probation departments across the state for those newly convicted of a prohibiting offense.

⁴For an abundance of examples of the product of these enforcement actions, see [California Department of Justice \(2015, 2016, 2017, 2018, 2019\)](#).

⁵While [Lott and Mustard \(1997\)](#) find the opposite, their results have been convincingly disputed on methodological grounds ([Aneja, Donohue III and Zhang, 2011](#); [National Research Council, 2005](#)).

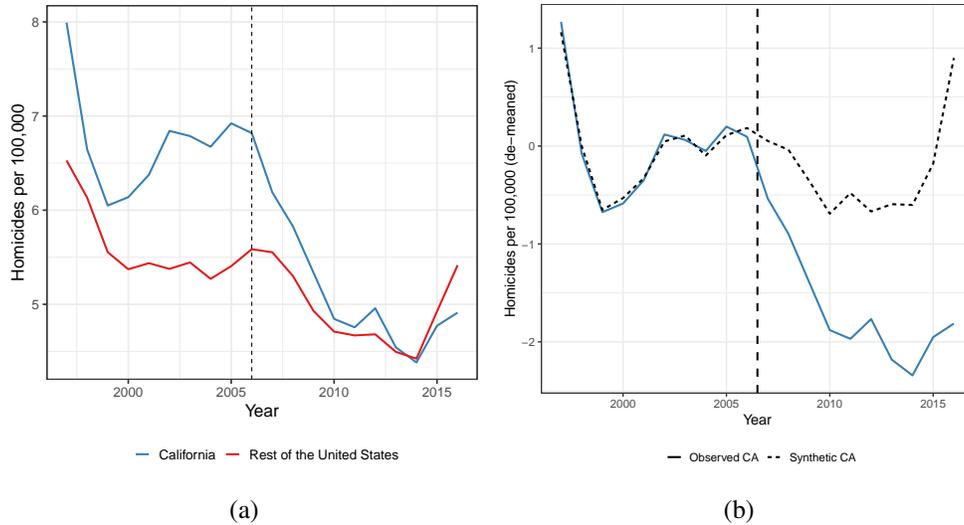


Fig 1: Annual homicide rate per 100,000, 1997 through 2016. The dotted line is 2006, the year the APPS program was launched. (a) California and the rest of the United States. (b) California and Synthetic California, matched based on the pre-intervention years 1997-2006, de-measured

gun-related deaths (Rudolph et al., 2015; Castillo-Carniglia et al., 2019). Methodologically, many of these studies also use the synthetic control method, which we consider next (Rudolph et al., 2015; Kagawa et al., 2018; Castillo-Carniglia et al., 2019; Donohue, Aneja and Weber, 2019; McCourt et al., 2020; Pear et al., 2022).

Several qualitative studies complement these quantitative estimates. Wintemute et al. (2019) present a series of California case studies of individuals subjected to gun law restraining orders, where the APPS was central in identifying associated firearms. In addition, case studies listed in recent years of the California Department of Justice annual APPS report indicate that BOF enforcement teams are often directly involved in carrying out these orders (California Department of Justice, 2018, 2019). In the case studies documented by Wintemute and colleagues, the individuals disarmed as a result of these orders all made specific and highly alarming threats against intimate partners, coworkers, educational institutions, and in some instances health and mental health providers. Among the cases where the individual in question threatened violence, none had committed a violent act at the time of publication. See also Swanson et al. (2017), who find that such restraining orders may also be effective in reducing gun suicide.

Finally, there is currently a randomized control trial of the APPS enforcement program in the field (see Wintemute et al., 2017, for the study protocol). The study began in 2015 to randomly assign regions within California for earlier or later enforcement actions among regions without well-established working relationships with the state Department of Justice. To date, results from this evaluation have not been reported.

2.3. Synthetic control estimates. Figure 1a shows the annual homicide rate per 100,000 people in California versus the rest of the United States (a population-weighted average of the remaining 49 states); the dotted line indicates 2006, the year the APPS program was first launched (in December). The figure is based on data from the Federal Bureau of Investigation (FBI) Uniform Crime Reports (UCR). The figure shows several striking patterns. First, the

overall murder rate in the United States largely declined over this period, with a major increase in the final years of the panel. Second, the murder rate in California was substantially higher than the national average in the late 1990s and mid 2000s, with an average difference in murders per 100,000 of 1.1 over the ten-year period through 2006. This gap closes after 2006: the California murder rate falls and remains below the rate for the rest of the nation by 2014.

The methodological question is whether and how we can use this panel data structure to estimate the impact of APPS on the murder rate. Furthermore, we would like to separately estimate effects on *gun related* and *non-gun related* homicides. The APPS program should primarily impact gun-related homicides by removing firearms from a high-risk population, though substitution effects may lead to changes in non-gun-related homicides. Therefore our methodological goal also includes estimating effects on *multiple outcomes* simultaneously.

Unfortunately, relatively simple approaches are unlikely to perform well here. For example, the workhorse “difference-in-differences” (DiD) regression model assumes *parallel trends*: in the absence of the intervention, the differences between the murder rates in California and the rest of the country are constant over time. This is clearly violated in Figure 1a, with much wider differences in the mid 2000s than in the late 1990s.

An alternative to DiD for settings such as this where the parallel trends assumption is implausible is the *synthetic control method* (SCM [Abadie, Diamond and Hainmueller, 2010](#)). This procedure finds a weighted combination of other states, *Synthetic California*, that approximates the pre-APPS murder rate for California. Here, we use a variant of SCM that first subtracts out the pre-APPS homicide rate in each state to isolate trends, then finds a synthetic control that approximates the pre-APPS murder rate trends for California ([Doudchenko and Imbens, 2016](#); [Ferman and Pinto, 2021](#); [Ben-Michael, Feller and Rothstein, 2021](#)). See Appendix A.3 for further details on constructing the synthetic control.

Figure 1b shows the de-meaned California murder rate (the solid line) alongside the de-meaned murder rate for our synthetic comparison group of states (the dashed line). The synthetic control has excellent pre-treatment fit: the two series are very close for the pre-treatment period of 1997 through 2006. Seven states have non-zero weight and contribute to the synthetic control: Ohio is by far the most important donor state (weight = 0.65), followed by Hawaii (0.1), Mississippi (0.09), Louisiana (0.06), Arkansas (0.06), Alaska (0.03) and Idaho (0.01). After 2006, California’s murder rate is markedly lower than the rate for synthetic California, with an average difference of -1.41 murders per 100,000. This difference widens over the ten-year period following the implementation of APPS.

While a promising initial analysis, several open questions remain. First, we would like to be able to estimate the effects on both gun related and non-gun related homicides in order to evaluate our hypothesis that the APPS program should have a larger impact on gun related homicides. Second, the methodology for panel data with multiple outcomes is underdeveloped, and it is unclear how this SCM approach can be applied more generally. Relatedly, the number of homicides per state are inherently count outcomes, a fact ignored in existing SCM analyses but one that we would like to incorporate explicitly. Finally, quantifying uncertainty with synthetic controls is challenging. The two most common approaches of placebo and conformal inference ([Abadie, Diamond and Hainmueller, 2010](#); [Chernozhukov, Wüthrich and Zhu, 2021](#)) rest on versions of exchangeability either over units or over time. These assumptions are difficult to justify in our setting, and so it is unclear whether the large differences in Figure 1b is driven by true effects or statistical noise. Moreover, the uncertainty intervals from conformal inference approaches are constant over time even though, intuitively, we anticipate far greater uncertainty for the counterfactual murder rate in 2017 than immediately after the implementation of APPS in 2007. This adds more difficulty in understanding whether the effect is truly increasing over time as we see in Figure 1b. These concerns motivate our alternative approach, which we turn to next.

3. Setup and Background.

3.1. *Problem setup.* We now describe the standard panel data setting with $i = 1, \dots, N$ units observed for $t = 1, \dots, T$ time periods. In our application, we observe $N = 50$ states for $T = 20$ years. Let W_i be an indicator that unit i is treated at time $T_0 < T$ where units with $W_i = 0$ never receive the treatment. While our setup is more general, we restrict our attention to the case where a single unit receives treatment, and follow the convention that this is the first one, $W_1 = 1$. The remaining $N_0 = N - 1$ units are untreated. In our application, the treatment of interest is the APPS program, treated at time $T_0 = 10$; California is the treated unit and the remaining states are possible controls. Note that while only California implement a program like APPS, as we discuss above, other states have enacted various gun-related policies. Thus, the un-treated condition $W_1 = 0$ corresponds to the policy environment California would have experienced in the absence of APPS.

We describe our problem of interest in the potential outcomes framework. We assume no interference between units and stable treatment, which allows us to write the potential outcomes for unit i in time t under control and treatment as $Y_{it}(0)$ and $Y_{it}(1)$ respectively.⁶ Our primary outcome of interest is the overall murder rate; we consider multiple outcomes in Section 4.3.3. For ease of exposition, we initially ignore the presence of background covariates from our analyses; we discuss extensions to incorporate them in Section 4.3.1.

Since the first unit is treated at time T_0 , the observed outcomes Y_{it} are therefore:

$$Y_{it} = \begin{cases} Y_{it}(0) & \text{if } W_i = 0 \text{ or } t \leq T_0 \\ Y_{it}(1) & \text{if } W_i = 1 \text{ and } t > T_0 \end{cases}.$$

The goal is to estimate the causal effect for the treated unit at each post-treatment time $t > T_0$, $\tau_t = Y_{1t}(1) - Y_{1t}(0) = Y_{1t} - Y_{1t}(0)$, as well as the average treatment effect on the treated unit after time T_0 ,

$$\tau = \frac{1}{T - T_0} \sum_{t=T_0+1}^T \tau_t.$$

Since we observe the treated outcome $Y_{1t} = Y_{1t}(1)$ for the treated unit at $t > T_0$, the challenge is to impute the missing potential outcome for the treated unit at time $t > T_0$. Note that we adopt a finite sample causal inference framework here: our goal is to estimate the impact for California, rather than for a hypothetical population (see [Imbens and Rubin, 2015](#)).

3.2. *Bayesian causal inference.* We adopt a Bayesian approach to estimating causal quantities. There are many subtle issues that arise in Bayesian causal inference that are not central to our discussion. For instance, identification in the Bayesian approach is conceptually distinct from identification in non-Bayesian causal inference: if the prior distribution is proper then the posterior distribution will also be proper, regardless of whether the parameters in the likelihood are fully or partially identified (see [Gustafson, 2010](#); [Imbens and Rubin, 2015](#)). See [Ding and Li \(2018\)](#) and [Oganisian and Roy \(2020\)](#) for additional background. [Pang, Liu and Xu \(2020\)](#) and [Menchetti and Bojinov \(2020\)](#) give overviews of Bayesian causal inference for panel or time series data.

At a high level, Bayesian causal inference views missing potential outcomes as unobserved random variables to be imputed ([Rubin, 1978](#); [Ding and Li, 2018](#)). We therefore break the

⁶While this is not always the case, in our application, it is reasonable to assume that any state could launch an APPS-style program at any point. Thus, both potential outcomes could reasonably exist for all units and all time periods.

problem into two parts: first, estimate a Bayesian model using the observed data; second, use this model to obtain the posterior (predictive) distribution for the missing potential outcomes. Specifically, let $\mathbf{Y}^{\text{mis}} = \{Y_{it}(0)\}_{i=1, t > T_0} \in \mathbb{R}^{T-T_0}$ be the missing potential outcomes for the treated unit after T_0 , and let $\mathbf{O}^{\text{obs}} = (\{Y_{it}(0)\}_{i > 1}, \{Y_{it}(0)\}_{i=1, t \leq T_0}, \{Y_{it}(1)\}_{i=1, t > T_0}, \mathbf{W})$ be the set of observed data, including all control outcomes, the treated unit’s pre-treatment outcomes, the treated unit’s observed outcomes post-treatment, and the set of treatment assignments. We introduce auxiliary covariates in Section 4.3.1.

The goal is to impute the missing potential outcomes, \mathbf{Y}^{mis} , by drawing from the posterior predictive distribution, $\mathbb{P}(\mathbf{Y}^{\text{mis}} | \theta, \mathbf{O}^{\text{obs}})$, governed by a model parameter θ . We can then estimate $\hat{\tau}_t = Y_{1t} - Y_{1t}^*(0)$ at post-treatment time $t > T_0$, where $Y_{1t}^*(0)$ is a draw from the posterior predictive distribution for $Y_{1t}(0)$ for unit 1 at time t . To do so, we follow [Richardson, Evans and Robins \(2011\)](#) and use the fact that:

$$\mathbb{P}\{\mathbf{Y}^{\text{mis}}, \theta | \mathbf{O}^{\text{obs}}\} = \mathbb{P}\{\theta | \mathbf{O}^{\text{obs}}\} \cdot \mathbb{P}\{\mathbf{Y}^{\text{mis}} | \theta, \mathbf{O}^{\text{obs}}\},$$

where $\mathbb{P}\{\mathbf{Y}^{\text{mis}}, \theta | \mathbf{O}^{\text{obs}}\}$ is the posterior distribution of our model parameters and the missing potential outcomes given observed data. Now, to obtain the target posterior predictive distribution, we first specify a prior distribution $\mathbb{P}(\theta)$ and use Bayes rule to obtain the posterior distribution $\mathbb{P}(\theta | \mathbf{O}^{\text{obs}})$. With this posterior over the model parameters, we can repeatedly draw samples of \mathbf{Y}^{mis} , yielding the joint posterior on the left-hand side of the above expression. Obtaining the target posterior predictive distribution is then a matter of rearranging the above expression to isolate $\mathbb{P}(\mathbf{Y}^{\text{mis}} | \mathbf{O}^{\text{obs}})$, marginalizing over θ . This approach rests entirely on the form of the model and the estimated model parameters. See [Pang, Liu and Xu \(2020\)](#) for an alternative framing that explicitly states sufficient conditions in the form of identifying assumptions, and [Miratrix \(2020\)](#) for a discussion of model-based uncertainty in a similar setting.

Formally, we assume that units and time periods are exchangeable conditional on the model parameters θ , which have a prior distribution $p(\theta)$. Following [Ding and Li \(2018\)](#), we partition the parameter space into θ^m , which governs the marginal potential outcome distribution, and θ^a , which governs the association between $Y_{it}(0)$ and $Y_{it}(1)$. In principle, we could treat θ^a as a sensitivity parameter; for the purposes of this paper, we assume that the potential outcomes are conditionally independent given data and model parameters, allowing us to ignore θ^a .⁷ The posterior for θ^m is then:

$$\mathbb{P}\{\theta^m | \mathbf{O}^{\text{obs}}\} \propto p(\theta^m) \prod_{(i,t) \in \mathcal{C}} \mathbb{P}(Y_{it}(0) | \theta^m, W_{it}),$$

where \mathcal{C} denotes the “control” unit and time periods: $t = 1, \dots, T_0$ for $i = 1$ and $t = 1, \dots, T$ for other units. Let $Y_{1t}^*(0)$ denote the posterior predictive distribution over the missing potential outcomes, $Y_{1t}^*(0) = \mathbb{P}\{\mathbf{Y}^{\text{mis}} | \theta, \mathbf{O}^{\text{obs}}\}$. The estimated treatment effect at time t is then the observed treated potential outcome minus the predicted missing potential outcome at post-treatment time $t > T_0$, $\hat{\tau}_t = Y_{1t} - Y_{1t}^*(0)$, where repeatedly drawing $Y_{1t}^*(0)$ from the posterior predictive distribution propagates the uncertainty.

A key consequence of this setup is that we never use post-treatment data from the treated unit in our modeling (or, equivalently, we have an arbitrarily flexible model for the observed treated data). Rather, our focus is on modeling the observed control potential outcomes. And by imputing $Y_{it}^*(0)$ from the posterior *predictive* distribution, our inference is inherently finite

⁷It is tempting to sidestep this issue entirely and view this framework as solely focusing on the control potential outcomes. The challenge is that a (formal) Bayesian approach conditions on all data and parameters. Thus, this additional restriction makes explicit the fact that we are not incorporating any information from the observed $Y_{it}(1)$. See [Ding and Li \(2018\)](#) for further discussion.

sample, allowing us to directly reason about the effect of APPS on California’s murder rate (Imbens and Rubin, 2015).

Finally, while the setup here is quite general, in practice we restrict ourselves to a structure where the control potential outcomes are a *model component* plus additive noise.

ASSUMPTION 1 (Additive, separable error structure). The control potential outcomes, $Y_{it}(0)$, are generated as:

$$Y_{it}(0) = \text{model}_{it}(\theta) + \varepsilon_{it} \quad \mathbb{E}[\varepsilon_{it}] = 0$$

where $\text{model}_{it}(\theta)$ is the outcome model for unit i at time t , governed by parameters θ ; and ε_{it} is mean-zero noise.

This assumption, sometimes referred to *strict exogeneity* in the panel data setting (see, e.g., Imai and Kim, 2019), places no restrictions on the model terms and θ . This allows for *non-parametric* modelling of the outcomes where θ is potentially infinitely dimensional — but does not allow treatment assignment to depend on the errors. In other words, given the model and the possibly infinitely many parameters θ , the error distribution for the treated unit’s (control) potential outcomes are exchangeable with other units’ errors. In our analysis of the APPS program, one implication of Assumption 1 is that California’s decision to implement APPS must not be in response to an *idiosyncratic shock* to the homicide rate. This is consistent with the history of the program (see, e.g., California Department of Justice, 2015), which largely came about due to changes in IT infrastructure and an entrepreneurial California Attorney General. As we discuss in Section 7, we can relax Assumption 1 within the fully Bayesian workflow, for instance, via sensitivity analysis (Franks, D’Amour and Feller, 2019).

3.3. *Related methodological research.* Our paper bridges several robust literatures. First, we build on the many existing methods for estimating causal effects for a single treated unit with panel data; see Samartsidis et al. (2019) for a recent review. Within this, there are several important threads. Most directly relevant are recent papers on Bayesian implementations of (or alternatives to) the synthetic control method, including Tuomaala (2019); Kim, Lee and Gupta (2020); Pang, Liu and Xu (2020); Pinkney (2021). Of particular interest is Brodersen et al. (2015), who propose a Bayesian structural time series model for estimating causal effects, focused on a single treated series; see Menchetti and Bojinov (2020) for a recent extension. We explore connections between Brodersen et al. (2015) and our proposed method in Appendix A.4. There are also several recent papers that directly estimate factor models for causal effects (e.g., Xu, 2017; Athey et al., 2021), as well as approaches that directly address multiple outcomes (Samartsidis et al., 2020). As we discuss below, we incorporate many of the novel ideas in these papers, including many of the prior choices. However, our proposed MTGP framework is typically more general, and, we believe, more conducive to estimating causal effects with panel data.

Next, there is a small but growing set of papers on the use of Gaussian Processes for estimating causal effects. To date, nearly all of these papers have focused on the cross-sectional setting (Alaa and van der Schaar, 2017; Schulam and Saria, 2017; Ray and van der Vaart, 2018; Huang, Chen and Liu, 2019; Branson et al., 2019; Witty et al., 2020; Ren et al., 2021). See Oganisian and Roy (2020) for a recent review.

We are aware of very few papers specifically using GPs to estimate causal effects in panel data settings, although, as we discuss in Section 5, many existing estimators can be written

as special cases of this approach.⁸ Two working papers are especially relevant for our work. The first is recent work from [Modi and Seljak \(2019\)](#), who also propose a GP approach for estimating causal effects in this setting. Their main proposal, however, is very different than ours, with a focus on GP for the frequency domain. Thus, we view our model as a useful complement to theirs. A second recent proposal is [Carlson \(2020\)](#), who also uses a GP approach but does not exploit the multitask structure. In independent work, [Chen et al. \(2022\)](#) also propose using a multitask GP for panel data. Unlike in our application, however, they consider the setting with many units divided into a treatment and control group and are focused on modeling so-called “parallel-ish” trends between these groups. In a different direction, [Antonelli and Beck \(2021\)](#) propose estimating effects under staggered adoption of treatment using a Bayesian vector autoregressive model that allows for spatial and time dependence, which is in some ways analogous to the multitask structure we propose. Finally, while not explicitly causal, of particular relevance to our approach is [Flaxman et al. \(2015\)](#), who propose a hierarchical GP model that is similar to what we discuss here.

4. Multitask Gaussian Processes for Causal Inference. We now give a high-level description of Multitask Gaussian Processes for causal effects with panel data, beginning with the simpler setting of single-task GPs. We discuss specific model implementation choices in [Appendix A](#). To fix terminology, we use *tasks* throughout to refer to units; thus, a single-task GP is for a single unit while multi-task GPs are for multiple units. This is in contrast to the original use (e.g., [Bonilla, Chai and Williams, 2008](#)) where tasks instead refers to outcomes.

4.1. Review: Single Task Gaussian Processes. We begin by reviewing the setup and properties of a Single-Task GP, which focuses solely on the treated unit; see [Williams and Rasmussen \(2006\)](#) for a textbook discussion. While the original GP setup is quite general, to fix ideas we initially focus on applying GPs to the *interrupted time series* or *horizontal regression* setting ([Athey et al., 2021](#)), in which we use the pre-treatment outcomes for California to forecast post-treatment outcomes in the absence of the intervention. Specifically, we initially model the (control) murder rate for California at time t , $Y_{1t}(0)$, as the sum of a model component f_{1t} and a mean-zero, independent noise component (as in [Assumption 1](#)):⁹

$$\begin{aligned} Y_{1t}(0) &= f_{1t} + \varepsilon_{1t} \\ \mathbf{f}_1 &\sim \mathcal{GP}(0, k_{\text{time}}) \\ \varepsilon_{1t} &\stackrel{\text{iid}}{\sim} N(0, \sigma^2) \end{aligned}$$

where $\mathcal{GP}(\cdot)$ is a Gaussian Process prior with corresponding *time kernel* k_{time} , $\mathbf{f}_1 = (f_{11}, \dots, f_{1T})$ is the vector of model components for the treated unit, and where we initially consider homoskedastic noise σ^2 , which we relax below.

The key idea is that the GP prior over \mathbf{f}_1 incorporates smoothness over time via $k_{\text{time}}(t, t')$, where similar values of t imply larger covariances. Here we use the squared exponential kernel:

$$k_{\text{time}}(t, t') = \exp\left(-\frac{|t - t'|^2}{2\rho}\right),$$

⁸See also [Karch, Brandmaier and Voelkle \(2018\)](#), who use GPs for latent growth curve modeling in psychometrics, and [Huang, Zhang and Schölkopf \(2015\)](#), who instead focus on causal discovery.

⁹For simplicity, we consider a de-meaned outcome series for now, allowing for a global intercept in [Section 4.3](#). Nonetheless, as we show in [Equation \(2\)](#), the posterior mean for a future data point can be nonzero even when \mathbf{f}_1 has mean zero; see [Rasmussen \(2004\)](#).

where ρ is the length scale.

The choice of kernel function plays an important role in defining the behavior of the model. For the family of kernels we consider, the relation between the time difference $|t - t'|$ and the kernel value $k_{\text{time}}(t, t')$ is central in determining how we expect the model to behave over time. If $k_{\text{time}}(t, t')$ is large even for far-apart time periods, then the model components $f_{1t}, f_{1t'}$ are assumed to be highly correlated; this corresponds to an assumption that the underlying model is very smooth over time. Conversely, if $k_{\text{time}}(t, t')$ is low even for close time periods, then the model components are assumed to be close to independent, which corresponds to assuming that the model varies strongly over time. The length scale hyper-parameter ρ explicitly controls the level of smoothness in the model components. Many other kernels are possible and appropriate for other settings, such as periodic kernels; again see [Williams and Rasmussen \(2006\)](#) for a textbook discussion.

This GP model implies that the control potential outcomes follow a multivariate Normal distribution, so the conditional distributions have a convenient form. Specifically, as in Section 3.2, we can partition the vector of control potential outcomes for the treated unit into the observed and missing components: $\mathbf{Y}_1^{\text{obs}} = (Y_{11}^{\text{obs}}, Y_{12}^{\text{obs}}, \dots, Y_{1T_0}^{\text{obs}})$, and $\mathbf{Y}_1^{\text{mis}} = (Y_{1T_0+1}(0), Y_{1T_0+2}(0), \dots, Y_{1T}(0))$. Next, we define $\mathbf{K}_{\text{time}} \in \mathbb{R}^{T \times T}$ as the *time kernel matrix*, where $\mathbf{K}_{\text{time}, tt'} = k_{\text{time}}(t, t')$. Suppressing the superscript *time* to reduce clutter, we can then partition this multivariate Normal model as:

$$(1) \quad \begin{pmatrix} \mathbf{Y}_1^{\text{mis}} \\ \mathbf{Y}_1^{\text{obs}} \end{pmatrix} \sim \text{MVN} \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \mathbf{K}_{\text{mis}} & \mathbf{K}_{\text{mis,obs}} \\ \mathbf{K}_{\text{obs,mis}} & \mathbf{K}_{\text{obs}} \end{pmatrix} + \sigma^2 \mathbb{I} \right],$$

where $\mathbf{K}_{\text{obs}} \in \mathbb{R}^{T_0 \times T_0}$ is the time kernel matrix for observed points, $\mathbf{K}_{\text{mis}} \in \mathbb{R}^{(T-T_0) \times (T-T_0)}$ is the time kernel matrix for unobserved points, and $\mathbf{K}_{\text{obs,mis}} \in \mathbb{R}^{T_0 \times (T-T_0)}$ is a rectangular matrix of time kernel evaluations for pre- and post-treatment times, whose t, t' element is the kernel value $k_{\text{time}}(t, T_0 + t')$.

We can then write the the posterior predictive distribution for California's outcomes at future time points, $t > T_0$, as:

$$\begin{aligned} (2) \quad \mathbf{Y}_1^{\text{mis}} | \mathbf{Y}_1^{\text{obs}} &\sim \text{MVN}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \\ \boldsymbol{\mu} &= \mathbf{K}_{\text{mis,obs}} (\mathbf{K}_{\text{obs}} + \sigma^2 \mathbb{I})^{-1} \mathbf{Y}_1^{\text{obs}} \\ (3) \quad \boldsymbol{\Sigma} &= (\mathbf{K}_{\text{mis}} + \sigma^2 \mathbb{I}) - \mathbf{K}_{\text{mis,obs}} (\mathbf{K}_{\text{obs}} + \sigma^2 \mathbb{I})^{-1} \mathbf{K}_{\text{obs,mis}}. \end{aligned}$$

As we discuss in Section 5, the posterior mean is a linear combination of the pre-treatment observations, $\mathbf{Y}_1^{\text{obs}}$.

4.2. Multitask Gaussian Processes. We now turn to modeling the murder rate for all states jointly, rather than modeling California's outcome series alone. The primary change is that each observation is now a unit-time pair (i, t) , so the kernel measuring similarity between points is doubly-indexed: $k((i, t), (i', t'))$. The *multitask* GP framework, originally developed for the setting with a single unit but many outcomes ([Goovaerts et al., 1997](#); [Bonilla, Chai and Williams, 2008](#)), is a natural approach for capturing similarity across multiple dimensions. As in the single-task GP setting and Assumption 1, we assume that the control potential outcomes consist of a model component plus additive (Normal) noise:

$$\begin{aligned} Y_{it}(0) &= f_{it} + \varepsilon_{it} \\ \mathbf{f} &\sim \mathcal{GP}(0, k) \\ \varepsilon_{it} &\stackrel{\text{iid}}{\sim} N(0, \sigma^2), \end{aligned}$$

where $\mathbf{f} = (f_{11}, \dots, f_{NT})$ is the vector of model components for all units and time periods. We again begin with the simplified setting in which the outcome series is de-measured by a global intercept; we incorporate a mean model in Section 4.3.

While this framework is quite general, we focus on two common simplifications: allowing for separable unit and time covariances, and imposing a low-rank structure on the unit covariance.

4.2.1. Separable unit and time kernels. The first key idea is to decompose the GP kernel into separate kernels that capture similarity across time (the “time covariance”) and similarity across units (the “unit covariance”):

$$(4) \quad k((i, t), (i', t')) = k_{\text{unit}}(i, i') \times k_{\text{time}}(t, t'),$$

where k_{unit} defines similarity between the N units and k_{time} defines similarity between the T time periods. This is a strong restriction that implies that the covariance across units, k_{unit} , is constant in time. (Conversely, this implies that the time covariance k_{time} is constant across units.) We discuss some approaches for relaxing this restriction in Section 7.

4.2.2. Low-rank unit covariance. Even with the separable kernel restriction, the resulting model is overparameterized. In particular, the similarity between the T time points, $k_{\text{time}}(t, t')$, has a natural parameterization in terms of the difference in time, $|t - t'|$, which is a scalar. By contrast, the similarity between the N units, $k_{\text{unit}}(i, i')$, involves $N(N - 1)/2$ comparisons. Following [Goovaerts et al. \(1997\)](#) and [Bonilla, Chai and Williams \(2008\)](#), we simplify the problem by using the *intrinsic coregionalization model* (ICM), which imposes a low-rank assumption on the unit covariance k_{unit} (see also [Flaxman et al., 2015](#), Section 4.2).

Specifically, we assume that the unit kernel matrix $\mathbf{K}_{\text{unit}} \in \mathbb{R}^{N \times N}$, where $\mathbf{K}_{\text{unit}, i, i'} = k_{\text{unit}}(i, i')$, has the form $\mathbf{K}_{\text{unit}} = \beta\beta'$ for some matrix $\beta \in \mathbb{R}^{N \times J}$, where J is the rank of the unit kernel matrix, and where we place a hyperprior on β and infer it from the data.¹⁰ While initially proposed in terms of restrictions on the matrices, we can equivalently write the ICM model as an assumption that the outcome model for each of the $i = 1, \dots, N$ units is a unit-specific linear combination of J latent draws from a Gaussian process, i.e.,

$$(5) \quad f_{it} = \sum_{j=1}^J \beta_{ij} u_{jt},$$

where each latent \mathbf{u}_j is itself a GP with kernel k_{time} , and where each unit has its own set of unit-specific factor loadings $\beta_i = (\beta_{i1}, \dots, \beta_{iJ})$. Thus, the underlying model components $\mathbf{u}_1, \dots, \mathbf{u}_J$ are shared across units and induce a dependence across them. The number of shared latent functions J corresponds to the rank of the unit kernel matrix \mathbf{K}_{unit} . This is a key hyper-parameter governing the overall complexity of the model; a higher rank J will lead to a more flexible model that potentially has the risk of overfitting. We discuss choosing the rank J via Bayesian model criticism in Section 6.1.¹¹ The resulting model is then:

$$Y_{it}(0) = f_{it} + \varepsilon_{it}$$

$$f_{it} = \sum_{j=1}^J \beta_{ij} u_{jt}$$

¹⁰Equivalently, we can write this restriction as $k_{\text{unit}}(i, i') = \langle \beta_i, \beta_{i'} \rangle$, i.e., a linear kernel between weights for units i and j , inferred from the data.

¹¹In principle, we could include a prior over the effective rank by shrinking columns of \mathbf{K}_{unit} to zero, though this is computationally demanding; see [Cheng et al. \(2020\)](#).

$$\mathbf{u}_j \sim \mathcal{GP}(0, k_{\text{time}})$$

$$\beta_{ij} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$$

$$\varepsilon_{it} \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2).$$

We estimate this model using the probabilistic programming language, Stan (Stan Development Team, 2021). See Appendix A for hyperprior distributions and additional implementation details; many other hyperprior distributions are possible.

As above, we can re-write this model in matrix notation and can represent the decomposition in Equation (4) as the *Kronecker product* between the two covariances:

$$\mathbf{K} \equiv \mathbf{K}_{\text{unit}} \otimes \mathbf{K}_{\text{time}} \in \mathbb{R}^{NT \times NT},$$

where any one entry of \mathbf{K} can be written as the product of the unit and time covariance, as in Equation (4).

We can similarly divide this (admittedly unwieldy) kernel matrix into a matrix for observed unit-time pairs, $\mathbf{K}_{\text{obs}} \in \mathbb{R}^{((n-1)T+T_0) \times ((n-1)T+T_0)}$, a matrix for the unobserved counterfactual post-treatment outcomes for the treated unit $\mathbf{K}_{\text{mis}} \in \mathbb{R}^{(T-T_0) \times (T-T_0)}$, and a rectangular matrix for the kernel between the outcomes for the observed unit-times and the post-treatment outcomes for the treated unit $\mathbf{K}_{\text{obs,mis}} \in \mathbb{R}^{((n-1)T+T_0) \times (T-T_0)}$. With this setup, the joint distribution of the missing post-treatment outcomes for the treated unit and the observed untreated outcomes are multivariate Normal:

$$\begin{pmatrix} \mathbf{Y}^{\text{mis}} \\ \mathbf{Y}^{\text{obs}} \end{pmatrix} \sim \text{MVN} \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \mathbf{K}_{\text{mis}} & \mathbf{K}_{\text{mis,obs}} \\ \mathbf{K}_{\text{obs,mis}} & \mathbf{K}_{\text{obs}} \end{pmatrix} + \sigma^2 \mathbb{I} \right].$$

The resulting conditional mean and variance have the same form as in Equations (2) and (3).

As we discuss in Section 5 below, there are two important special cases of this model, which correspond to setting either the time or unit covariances to be the identity. Setting the *time* covariance matrix to the identity, $\mathbf{K}_{\text{time}} = \mathbb{I}_T$, corresponds to a model with no smoothness in the underlying latent factors, \mathbf{u} , allowing the function value to change arbitrarily between time points. The resulting approach is a (Bayesian) linear factor model (e.g., Xu, 2017; Athey et al., 2021; Samartsidis et al., 2020). In Section 5 below, we further connect this to “vertical regression,” which finds a set of weights over units that optimize some imbalance criterion (Doudchenko and Imbens, 2016). On the other hand, setting the *unit* covariance matrix to the identity, $\mathbf{K}_{\text{unit}} = \mathbb{I}_N$, corresponds to a model in which there is no correlation in the underlying process for each unit, and each model component is an independent GP. The model therefore reduces back to the single-task GP discussed in Section 4.1 above. Finally, while we focus on the ICM model with separate time and unit covariances, there is a large literature on more elaborate generalizations, such as the *semiparametric latent factor model* and the *linear model of coregionalization*; see Álvarez (2017) for a review.

4.2.3. Computational considerations. An important practical constraint in applying multi-task Gaussian processes is the prohibitive computational burden incurred when naively taking samples. Specifically, the Cholesky decomposition of the covariance matrix defined by $\mathbf{K}_{\text{time}} \otimes \mathbf{K}_{\text{unit}}$ has computational complexity cubic with respect to the product of the number of units and time points, $\mathcal{O}((TN)^3)$. Fortunately, we can take advantage of the structure of the covariance matrix to dramatically reduce this complexity to $\mathcal{O}(T^3 + NJ)$ using properties of Kronecker products and low-rank approximations to each constituent matrix. Let \mathbf{C}_{time} and \mathbf{C}_{unit} denote the decomposition of the time and unit matrices, \mathbf{K}_{time} and \mathbf{K}_{unit} , such that $\mathbf{K} = \mathbf{C}\mathbf{C}^\top$, and let \mathbf{C}_{\otimes} be the decomposition of the Kronecker product of \mathbf{K}_{time} and \mathbf{K}_{unit} . From the mixed product property of Kronecker products, we have

$$\mathbf{C}_{\otimes} \mathbf{C}_{\otimes}^\top = (\mathbf{C}_{\text{time}} \otimes \mathbf{C}_{\text{unit}}) (\mathbf{C}_{\text{time}} \otimes \mathbf{C}_{\text{unit}})^\top,$$

which allows for the separate decomposition of each covariance. This is $\mathcal{O}(T^3)$ for the time covariance,¹² and $\mathcal{O}(NJ)$ for the unit covariance, which uses our low-rank representation $\mathbf{K}_{\text{unit}} = \beta\beta^\top$, $\beta \in \mathbb{R}^{N \times J}$.¹³

4.3. *Extensions.* Building on the broad GP literature, we can immediately extend the MTGP model above to better match our application. We focus on three main extensions here: incorporating a mean model; allowing for a non-Gaussian likelihood; and modeling multiple outcomes simultaneously.

4.3.1. *Incorporating a mean model.* Thus far, we have considered mean-zero Gaussian Processes for the structural component of the control potential outcomes. A natural extension is to instead allow for the GP to have a non-zero mean, m_{it} :¹⁴

$$\begin{aligned} Y_{it}(0) &= f_{it} + \varepsilon_{it} \\ \mathbf{f} &\sim \mathcal{GP}(\mathbf{m}, k) \\ \varepsilon_{it} &\sim N(0, \sigma^2), \end{aligned}$$

where $\mathbf{m} = (m_{11}, \dots, m_{NT})$ is a length $N \times T$ prior mean vector. Extending the formula in Equation (2), the posterior mean for a new control potential outcome is therefore:

$$\boldsymbol{\mu} = \mathbf{m} + \mathbf{K}_{\text{mis,obs}}(\mathbf{K}_{\text{obs}} + \sigma^2\mathbb{I})^{-1}(\mathbf{Y}^{\text{obs}} - \mathbf{m}).$$

We return to this representation in Section 5.1 below.

Explicitly incorporating a mean model is useful in a panel data setting because we can more directly connect the MTGP framework to common estimators, such as Comparative Interrupted Time Series and so-called fixed effects models (for a recent review, see [Samartsidis et al., 2019](#)). One natural model is to set m_{it} to have unit- and time-specific intercepts:

$$\begin{aligned} m_{it} &= \mu + \text{unit}_i + \text{time}_t \\ \text{unit}_i &\sim N(0, \sigma_{\text{unit}}^2) \\ \text{time} &\sim \mathcal{GP}(0, k_{\text{global}}) \end{aligned}$$

where μ is a global intercept, unit_i is a unit-specific intercept for unit i , and time_t is a time-specific intercept for time t . The unit-specific intercepts capture the fact that one state may have a higher baseline rate of homicides than another, even if they have similar trends. The time-specific intercepts capture the trend in the homicide rate across the United States. This allows us to estimate and account for underlying trends in the homicide rate over time. Here we place independent Normal priors on the unit-specific intercepts, and a global Gaussian Process prior on the time-specific intercepts. We could similarly allow for a parametric trend

¹²In cases where T is large, a number of methods can be used to improve computational efficiency further by reducing the complexity of taking the decomposition of \mathbf{K}_{time} (see [Solin and Särkkä, 2020](#); [Wilson, Dann and Nickisch, 2015](#); [Flaxman et al., 2015](#)).

¹³We can further reduce the computational overhead incurred for sampling from this distribution. Let $\mathbf{Z} \in \mathbb{R}^{T \times J}$, $\mathbf{Z}_{i,j} \sim \mathcal{N}(0, 1)$. Sampling by taking $\tilde{\mathbf{z}} = \mathbf{C}_{\text{time}} \otimes \mathbf{C}_{\text{unit}} \text{vec}(\mathbf{Z})$ is naively $\mathcal{O}(TJ)$. We can reduce this to $\mathcal{O}(T^2 + NJ)$ by using the “vectorization trick,” $(\mathbf{C}_{\text{time}} \otimes \mathbf{C}_{\text{unit}}) \text{vec}(\mathbf{Z}) = \text{vec}(\mathbf{C}_{\text{unit}} \mathbf{Z} \mathbf{C}_{\text{time}}^\top)$.

¹⁴The distinction between the mean function and the covariance matrix in the MTGP is somewhat artificial, and is typically chosen for clarity and to facilitate the choice of prior; see [Rasmussen \(2004, Ch. 2.7\)](#) and [Kana-gawa et al. \(2018\)](#). In this application, making the mean function explicit allows us to more directly connect the proposed approach to the panel data literature.

over time, such as a baseline linear model common in Interrupted Time Series approaches (see [Miratrix, 2020](#)).

Finally, we can extend the mean model to incorporate (time invariant) auxiliary covariates, $X_i \in \mathbb{R}^p$:¹⁵

$$m_{it} = \mu + \text{unit}_i + \text{time}_t + \eta' X_i,$$

where η is a vector of regression coefficients, that we can place a prior on. As with the unit-specific intercepts, we could instead include this as part of the GP kernel, for example by incorporating it into the unit covariance.

4.3.2. Non-Gaussian Likelihoods and heteroskedasticity. Another natural GP extension is to allow for non-Gaussian observation models, which is essential for well-calibrated uncertainty quantification. Analogous to using alternative link functions with generalized linear models ([Gelman et al., 2013](#), Ch. 21.3), we can model the control potential outcomes as:

$$Y_{it}(0) \sim g^{-1}(f_{it}),$$

where $g^{-1}(\cdot)$ is an appropriate link function and f_{it} is the latent GP factor. As above, we could also incorporate a mean model, $Y_{it}(0) \sim g^{-1}(m_{it} + f_{it})$. See [Naish-Guzman and Holden \(2008\)](#); [Hensman, Matthews and Ghahramani \(2015\)](#) for discussions of related computational issues.

We highlight two link functions here. First, the outcome in our application is the number of murders, which is more naturally modeled via a Poisson (or, perhaps, Negative Binomial) link. A Poisson model also captures the restriction that the variance in the murder rate increases with the mean murder rate. We use this as our primary model below. Alternatively, we could allow for a “robust” Gaussian Process by allowing for t -distributed rather than Normal errors; see [Jylänki, Vanhatalo and Vehtari \(2011\)](#). In our application, the estimates using t -distributed errors are nearly identical to those using Normal errors.

Finally, an important practical feature of our application is that populations vary considerably across states. Thus, we expect that the underlying variability in the murder rate will be much larger in, say Vermont or Wyoming than in California or Texas. With a Gaussian likelihood, we can parametrize this as:

$$Y_{it}(0) \sim N\left(f_{it}, \frac{\sigma^2}{N_{it}}\right)$$

where N_{it} is the population for state i at time t . We can similarly allow for a population offset in a Poisson model.¹⁶ Allowing for heteroskedasticity by size is an important departure from many Frequentist panel data methods used in this setting; see [Samartsidis et al. \(2019\)](#).

4.3.3. Multiple outcomes. We can extend the MTGP approach to allow for multiple correlated outcomes, following the large literature on multi-output GPs ([Álvarez, 2017](#)); see also [Samartsidis et al. \(2020\)](#) for a non-GP setting. This is particularly important for our application because we are interested in assessing the impact of APPS on both gun- and non-gun-related homicides.

To do so, we can write each data point as a unit-time-outcome triple, (i, t, ℓ) , with control potential outcome, $Y_{it\ell}(0)$. While many models are possible, we focus on a simple version

¹⁵With additional restrictions (e.g., “exogeneity”), we can extend this to include time-varying covariates as well. See [Imai and Kim \(2019\)](#) for discussion.

¹⁶More broadly, we could incorporate much more complex noise models, such as an additional GP on the noise term (e.g., [Naish-Guzman and Holden, 2008](#); [Hensman, Matthews and Ghahramani, 2015](#)). This is a challenging model to fit given our limited sample size.

with a shared factor structure across outcomes, which we can write as a multivariate Normal observation model with common covariance across units and time:

$$\mathbf{Y}_{it}(0) \sim \text{MVN}(\mathbf{f}_{it}, \Sigma),$$

where $\mathbf{f}_{it} = (f_{it1}, \dots, f_{itL})$ and $\mathbf{Y}_{it}(0) = (Y_{it1}(0), \dots, Y_{itL}(0))$ for outcomes $\ell = 1, \dots, L$. Equivalently, we can write this as an MTGP where the resulting kernel decomposes into separate unit, time, and outcome kernels:

$$k((i, t, \ell), (i', t', \ell')) = k_{\text{unit}}(i, i') \times k_{\text{time}}(t, t') \times k_{\text{outcome}}(\ell, \ell').$$

Similarly, we can write this as a Kronecker structured overall kernel $\mathbf{K} = \mathbf{K}_{\text{unit}} \otimes \mathbf{K}_{\text{time}} \otimes \mathbf{K}_{\text{outcome}}$, as in [Flaxman et al. \(2015\)](#). Unlike the unit kernel $k_{\text{unit}}(\cdot, \cdot)$, which we restrict to be low-rank, the outcome kernel $k_{\text{outcome}}(\cdot, \cdot)$ can be full rank because the number of outcomes (2) is less than the number of units (50).

5. Guarantees and connections through a weighting representation. Our discussion thus far has been entirely Bayesian. Gaussian Processes, however, are fundamentally linked with (Frequentist) kernel ridge regression; a large literature exploits this connection to describe the Frequentist property of GPs; see [Kanagawa et al. \(2018\)](#) for a recent review. Here we adapt those results to our panel data setting. We first represent the MTGP estimate as a weighting estimator with separate unit and time weights, and give error bounds under fairly general conditions on the latent (noiseless) outcomes. We then describe several special cases to better illustrate this result.

5.1. Weighting representation. We now show that the MTGP estimate with separable unit and time kernels has a corresponding representation as a weighting estimator with separate unit and time weights. This setup differs from our fully Bayesian formulation above; to make this connection with weighting, we instead consider a Frequentist setup that assumes a known kernel. Specifically, conditional on hyperparameters, the MTGP posterior mean can be written as the solution to a constrained optimization problem that minimizes the worst-case mean-square error across the “noiseless” latent functions, f . This representation exploits the connection between Gaussian Processes and kernel ridge regression ([Kanagawa et al., 2018](#)); see [Hazlett and Xu \(2018\)](#) for additional discussion of kernel weighting methods for panel data.

To set up the problem, we again assume that (control) potential outcomes are the sum of structural and error components, $Y_{it}(0) = f_{it} + \varepsilon_{it}$, where ε_{it} is mean-zero noise. We assume that f is “well behaved” in the sense of belonging to a Reproducing Kernel Hilbert Space (RKHS); this setup is quite flexible and incorporates a wide range of function classes (see, e.g., [Wainwright, 2019](#)). We focus on the ICM kernel, $\mathbf{K} = \mathbf{K}_{\text{time}} \otimes \mathbf{K}_{\text{unit}}$, where, as in other Frequentist analyses, the kernel is assumed known.

We can then represent the MTGP as a weighting estimator. We consider two basic forms, corresponding to the weight-space and function-space interpretations of Gaussian processes ([Williams and Rasmussen \(2006\)](#), Ch. 2). We first show that the MTGP minimizes the worst-case imbalance in the noiseless latent functions f . We then show the equivalent formulation in terms of imbalance on the observed outcomes Y . We initially consider mean-zero GPs, and incorporate a prior mean function below.

Let $Y_{it}(0) = f_{it} + \varepsilon_{it}$, where f is a fixed function. Let \mathcal{H}_k be the RKHS for kernel k with Hilbert-Schmidt norm, $\|f\|_{\mathcal{H}_k}^2 = k(f, f)$. Then let f be contained in the unit ball of the reproducing Hilbert space, $\|f\|_{\mathcal{H}_k} \leq 1$, where $k((i, t), (i', t')) = k_{\text{unit}}(i, i') \times k_{\text{time}}(t, t')$, and ε_{it} are iid mean-zero random variables with observed variance $\sigma^2 > 0$. Further, model the time covariance \mathbf{K}_{time} with a squared exponential kernel with length scale ρ , and model the

unit covariance $\mathbf{K}_{\text{unit}} = \beta\beta^\top$. Finally, let $\gamma^{(i)} \in \mathbb{R}^N$ and $\lambda^{(t)} \in \mathbb{R}^T$ be the vectors of unit and time weights for target unit i at time t , and let γ and λ , respectively, be the set of weights across all targets. Then:

(a) The posterior (predictive) mean estimate for target observation $(1, t^*) \notin \mathcal{C}$, for $t^* > T_0$ is:

$$\hat{\mu}_{1t^*} = \sum_{(i,t) \in \mathcal{C}} \gamma_i^{(1)} \lambda_t^{(t^*)} Y_{it}.$$

The unit and time weights, γ and λ , have the following closed form:

$$(6) \quad \gamma^{(i)} \otimes \lambda^{(t)} = (\mathbf{K}_{\text{time}} \otimes \mathbf{K}_{\text{unit}} + \sigma^2 \mathbb{I})^{-1} (k_{\text{time}}(t, \cdot) \otimes k_{\text{unit}}(i, \cdot)).$$

(b) The unit and time weights also solve the following, equivalent optimization problems, in terms of (1) the weight-space formulation:

$$(7) \quad \min_{\gamma, \lambda} \sup_{f \in \mathcal{H}_k} \sum_{(i,t) \in \mathcal{C}} \left(f_{it} - \sum_{(i',t') \in \mathcal{C}} \gamma_{i'}^{(i)} \lambda_{t'}^{(t)} f_{i't'} \right)^2 + \sigma^2 \|\gamma\|_2^2 \|\lambda\|_2^2,$$

and (2) the function-space formulation:

$$(8) \quad \min_{\alpha, \xi} \sum_{(i,t) \in \mathcal{C}} \left(Y_{it} - \sum_{(i',t') \in \mathcal{C}} \underbrace{\alpha_{i'}^{(i)}}_{\gamma_{i'}^{(i)}} k_{\text{unit}}(i, i') \cdot \underbrace{\xi_{t'}^{(t)}}_{\lambda_{t'}^{(t)}} k_{\text{time}}(t, t') Y_{i't'} \right)^2 + \sigma^2 (\xi \otimes \alpha)^\top \mathbf{K} (\xi \otimes \alpha),$$

where $\mathbf{K} = \mathbf{K}_{\text{time}} \otimes \mathbf{K}_{\text{unit}}$ and where $\alpha^{(i)} \in \mathbb{R}^N$ and $\xi^{(t)} \in \mathbb{R}^T$ are the coefficient vectors for target unit i at time t , with $\gamma_{i'}^{(i)} = \alpha_{i'}^{(i)} k_{\text{unit}}(i, i')$ and $\lambda_{t'}^{(t)} = \xi_{t'}^{(t)} k_{\text{time}}(t, t')$, and with corresponding sets α and ξ .

(c) This estimate has the following out-of-training-sample estimation error for the structural component of the missing potential outcome f_{1t^*} ,

$$(9) \quad |f_{1t^*} - \hat{\mu}_{1t^*}| \leq \left[\underbrace{\left(k_{\text{time}}(t^*, \cdot)^\top \lambda^{(t^*)} \otimes k_{\text{unit}}(1, \cdot)^\top \gamma^{(1)} \right)}_{\text{error for } f_{1t^*}} + \underbrace{\sigma^2 \left\| \gamma^{(1)} \right\|_2^2 \left\| \lambda^{(t^*)} \right\|_2^2}_{\text{irreducible noise}} \right]^{1/2}.$$

Proposition 5.1 begins with the algebraic result that the posterior mean estimate of the MTGP can be written as a weighting estimator, and that these overall weights separate into unit- and time-weights. This separation does not hold in general, but, in our setting with an ICM kernel, follows immediately from the restriction that the unit and time covariances are separable (Bonilla, Chai and Williams, 2008). As a result, we can view our proposed MTGP approach as a doubly-weighted panel data estimator, similar to recent proposals from e.g. Ben-Michael, Feller and Rothstein (2021) and Arkhangelsky et al. (2019) that also impute the counterfactual as a double weighted average.¹⁷ Importantly, the MTGP implicitly estimates unit and time weights *simultaneously*, while previous proposals construct them separately.

¹⁷The weights we consider here can be negative, unlike some Frequentist weighting methods. Restricting the weights to obey a simplex constraint can be achieved by using the graph Laplacian and considering the Gaussian random field, an analogue of Gaussian processes (Zhu, 2005).

Proposition 5.1 then leverages Frequentist results on GPs to provide additional intuition for these weights (see Kanagawa et al., 2018). First, Equation (7) shows that the MTGP weights minimize the error on the noiseless outcome functions, subject to regularization on the weights. Importantly, this formulation does not rely on parametric assumptions, only that the noiseless outcomes, f_{it} , are relatively “well behaved” in the sense of belonging to an RKHS. Equation (8) shows that the weights can be seen as the solution to an equivalent optimization problem in terms of minimizing imbalance in the observed outcomes directly, but subject to regularization on the “coefficients” α and ξ , rather than on the weights themselves. Finally, building off this, Equation (9) shows that the error bound depends directly on the distance between the control observations, \mathcal{C} , and the unobserved point, $(1, t^*)$. Specifically, with a squared exponential kernel, $k_{\text{time}}(t^*, \cdot)$, the bound is increasing for $t^* > T_0$. See Fiedler, Scherer and Trimpe (2021) for recent generalizations of such bounds under model mis-specification.

5.1.1. *Using a prior mean model.* The results above are for the MTGP where the prior mean model m_{it} is zero. We can extend these results to include the prior mean model m_{it} from Section 4.3.1, recognizing that the distinction between modeling the mean and modeling the covariance can often be arbitrary (Williams and Rasmussen, 2006). The resulting posterior (predictive) mean estimate is

$$(10) \quad \hat{\mu}_{1t^*}^{\text{aug}} = \sum_{(i,t) \in \mathcal{C}} \gamma_i^{(1)} \lambda_t^{(t^*)} Y_{it} + \underbrace{\left(m_{1t^*} - \sum_{(i,t) \in \mathcal{C}} \gamma_i^{(1)} \lambda_t^{(t^*)} m_{it} \right)}_{\text{estimated bias}}$$

$$(11) \quad = m_{1t^*} + \underbrace{\sum_{(i,t) \in \mathcal{C}} \gamma_i^{(1)} \lambda_t^{(t^*)} (Y_{it} - m_{it})}_{\text{re-weighted residuals}}.$$

Equation (10) has a form similar to bias correction for matching (Rubin, 1973), which corrects the doubly-weighted average by an estimate of the difference in the outcome model m_{it} . Equation (11) instead has the form of augmented inverse propensity score weighting (Robins, Rotnitzky and Zhao, 1994), which re-weights the outcome model residuals. This echos several recent proposals that combine outcome modeling and weighting in the panel data setting (see, e.g. Ben-Michael, Feller and Rothstein, 2021; Ferman and Pinto, 2021; Arkhangelsky and Imbens, 2021), where here the “outcome model” is our prior for the control potential outcome and can include, e.g. fixed effects and time invariant auxiliary covariates as in Section 4.3.1.

5.2. *Special cases.* To help build intuition for the weighting representation, we consider two special cases: (1) the unit kernel \mathbf{K}_{unit} is the identity; and (2) the time kernel \mathbf{K}_{time} is the identity matrix. For simplicity, we take the prior mean model to be zero here, but could extend accordingly.

5.2.1. *Identity unit covariance.* First, let the unit covariance \mathbf{K}_{unit} be the identity matrix. We can view this as the setting where there is no correlation between the underlying processes for each unit, so each model component is an independent GP. In this setting, there is nothing to be learned from the comparison units, and the MTGP estimate only compares to the pre-treatment outcomes of the treated unit. In particular, following the setup in Proposition 5.1, if we set the unit covariance matrix \mathbf{K}_{unit} to be the identity, then the weights for post-treatment time period $t^* > T_0$ are

$$\lambda^{(t^*)} = (\mathbf{K}_{\text{time}} + \sigma \mathbb{I}_{T_0})^{-1} k_{\text{time}}(t^*, \cdot),$$

with posterior mean estimate:

$$\hat{\mu}_{1t^*}^{\text{horiz}} = \sum_{t=1}^{T_0} \lambda_t^{(t^*)} Y_{1t}.$$

This is a *horizontal regression* formulation (see [Athey et al., 2021](#)). As these weights ignore the outcomes for the comparison units, we can view this case as fitting N separate single-task GP models, one for each unit. Thus, when the unit covariance is the identity matrix, this problem ignores the outcomes for comparison units and reduces to the single-task GP for the treated unit discussed in Section 4.1.

5.2.2. Identity time kernel. Next, let the time kernel \mathbf{K}_{time} be the identity matrix. This corresponds to the setting where there is no smoothness at all in the underlying latent functions, so the function value can change wildly between time periods. In this case, it is impossible to use outcomes from different time periods to inform our estimates, and so only the unit weights remain. Again following the setup in Proposition 5.1, if the time covariance matrix \mathbf{K}_{time} is the identity, then the weights for the treated unit are:

$$\gamma^{(1)} = (\mathbf{K}_{\text{unit}} + \sigma \mathbb{I}_N)^{-1} k_{\text{unit}}(1, \cdot),$$

with posterior mean estimate:

$$\hat{\mu}_{1t^*}^{\text{vert}} = \sum_{i=2}^N \gamma_i^{(1)} Y_{it^*}.$$

This is a *vertical regression* setup (see [Doudchenko and Imbens, 2016](#)), and corresponds to the implicit weights of a linear factor model or interactive fixed effects model ([Gobillon and Magnac, 2016](#); [Xu, 2017](#); [Athey et al., 2021](#)). The estimator only uses the learned unit covariance to implicitly construct weights over comparison units.

6. Model diagnostics and estimated impact for APPS. We now use the MTGP framework to estimate the impact of APPS on homicides. We begin by using posterior predictive checks to guide the choice of model. We then estimate the impacts both on overall homicides and separately on gun- and non-gun-related homicides.

6.1. Posterior Predictive Checks and other model diagnostics. Diagnostics and other forms of model criticism are important components of the Bayesian workflow, helping researchers diagnose and compare the fit of the proposed models ([Gelman et al., 2013](#)). We begin with *posterior predictive checks* (PPCs). First, we choose a test statistic, $T(\text{data}, \theta)$, which reflects a relevant aspect of model fit and is a function of both data and model parameters. We then compare the posterior predictive distribution of this statistic, $T(\text{data}_{\text{rep}}, \theta)$, to the value of the test statistic for the observed data, $T(\text{data}_{\text{obs}}, \theta)$, over draws of θ .

PPCs play an especially important role in panel data settings like ours, where it is often difficult to decide between modeling approaches given limited data.¹⁸ We focus on two common measures of model fit: the overall pre-treatment fit and the pre-treatment imbalance at each time point. First, we evaluate the pre-treatment fit for California by comparing the posterior predictive distribution of the root mean squared error (RMSE),

$$T(\mathbf{Y}^*, \theta) = \frac{1}{T_0} \sqrt{\sum_{t=1}^{T_0} (Y_{1t}^* - f_{1t})^2},$$

¹⁸See [Liu, Wang and Xu \(2020\)](#) for a review of Frequentist approaches to model checks for panel data.

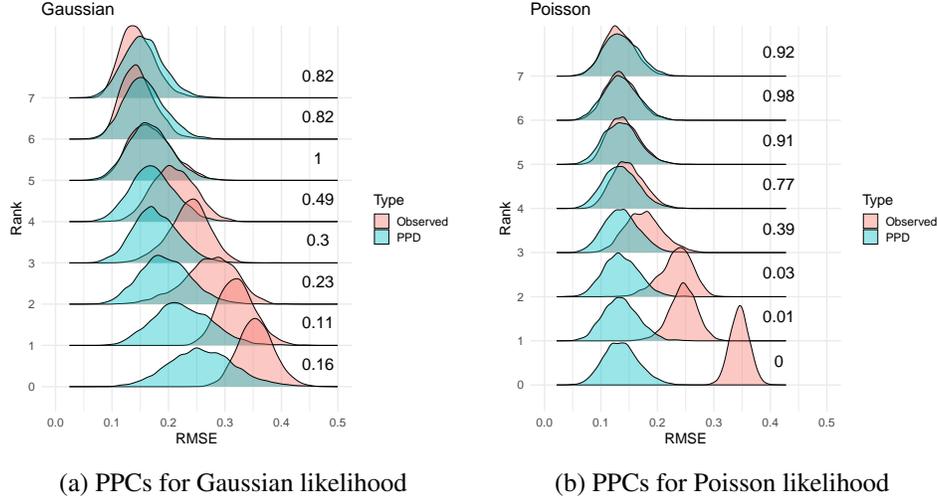


Fig 2: Posterior Predictive Checks (PPCs) for the pre-treatment RMSE for California. Numbers indicate posterior predictive p-value associated with each rank, i.e., $\mathbb{P}(T(\text{data}_{\text{rep}}, \theta) > T(\text{data}_{\text{obs}}, \theta) \mid \text{data})$.

to the observed RMSE, $T(\mathbf{Y}^{\text{obs}}, \theta) = \frac{1}{T_0} \sqrt{\sum_{t=1}^{T_0} (Y_{1t} - f_{1t})^2}$, where Y_{1t}^* is a posterior predictive draw of the overall murder rate for California at (pre-treatment) time t . Figure 2 shows the distributions for both Gaussian and Poisson observation models with unit kernel ranks $J = 0, \dots, 7$. For each combination, we also compute the posterior predictive p-value, the fraction of posterior predictive test statistics that are larger than the observed test statistic. For both Gaussian and Poisson models, the observed RMSE is larger than we would expect for ranks 0 through 3, with minimal discrepancies by rank 5 for both models. The Poisson model has slightly lower RMSE, suggesting that Poisson with rank 5 is a reasonable initial choice here.

Figure 3 shows the imbalance at each pre-treatment time point for the Poisson MTGP model with rank 5. The upper pane shows the observed data (in red) as well as posterior predictive draws from the model on the scale of the original outcome (homicides per 100,000). The lower pane instead shows the “gap” between the outcome and the posterior mean, $T(\text{data}, \theta) = Y_{1t} - \hat{\mu}_{1t}$.¹⁹ Both plots show little difference between the observed data and the posterior predictive distribution, suggesting reasonable model fit along the dimensions we investigate here.

Finally, the Appendix includes several additional model diagnostics. Appendix Figure C.3 shows the coverage of the 50% and 95% posterior predictive intervals for pre-treatment outcomes across all 50 states, computed as $\frac{1}{NT} \sum_{i,t} I(L_{it}^q < Y_{it} < U_{it}^q)$ where L_{it}^q and U_{it}^q are the lower and upper endpoints of the credible region with probability q for unit i at time t . The Poisson intervals have close to nominal coverage for ranks 4 and 5, but under-cover for lower ranks and over-cover for higher ranks. By contrast, the corresponding Gaussian intervals over-cover across all ranks, with especially poor coverage for the 50% intervals for higher ranks, leading further credence to the Poisson model over the Gaussian model. This difference is likely due to the fact that the sampling variance increases with the mean in the Poisson model but is constant in the mean and across units for the Gaussian model.

¹⁹This is the Bayesian analog to estimates of placebo impacts that are standard in Frequentist panel data studies, although we compare the observed estimates to the posterior predictive distribution rather than to zero.

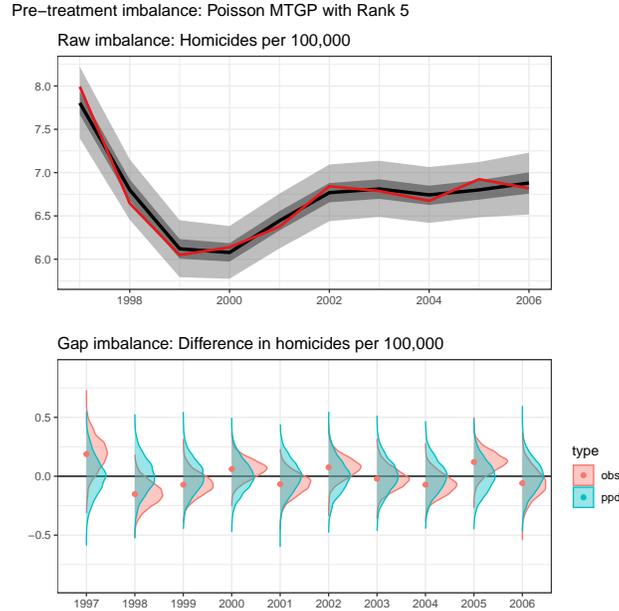


Fig 3: Posterior predictive checks for pre-treatment imbalance with the Poisson model with Rank 5. Observed data is in red. The upper panel shows the posterior predictive distribution and original data on the scale of the raw imbalance. The lower panel shows the corresponding “gap” estimates for each year, comparing the distribution of $Y_{1t} - \hat{\mu}_{1t}$ to the distribution of $Y_{1t}^* - \hat{\mu}_{1t}$ for times $t = 1, \dots, T_0$. These placebo estimates are overlapping, giving further support to good model fit.

In Appendix Figures C.4 and C.5 we show the inferred Gaussian Process for the time-specific intercept and the estimated factors from the rank 5 Poisson model next to the factor loadings for California. In Appendix Figures C.6 and C.7, we show the inferred time and unit weights for California in the first post-treatment period using a rank 5 Gaussian multitask Gaussian process without the global time series trend GP. Each of the time and unit weights are obtained by marginalizing over the non-relevant component (unit, or time, respectively) for each sample from the MCMC sampler.

6.2. *Impact of APPS on homicides.* We now turn to estimating the impact of APPS on homicides in California. We begin with the overall impact and then disaggregate by gun and non-gun-related homicides.

6.2.1. *Overall impact of APPS.* Based on our model diagnostics above, we focus on the MTGP estimated with rank $J = 5$, Poisson observation model, and unit- and time-specific intercepts. Figure 4 shows the estimated impact of APPS on homicide rates in California over time.²⁰ The upper pane shows the posterior predictive distribution from the MTGP, with

²⁰Appendix Figure C.11 shows the corresponding estimates after adjusting for auxiliary covariates and several measures of crime, all for 2005. The auxiliary covariates are: the prison incarceration rate, the age distribution (percent 0-17, 18-24, 25-44, 45-64, 65 and older), the percent Black, the unemployment rate, the poverty rate, and log median income. The crime measures, normalized by population, are: overall violent crime, rape, robbery, assault, property crime, burglary, larceny, and motor/vehicular crime. As the results are largely unchanged, we focus on the unadjusted estimate here.

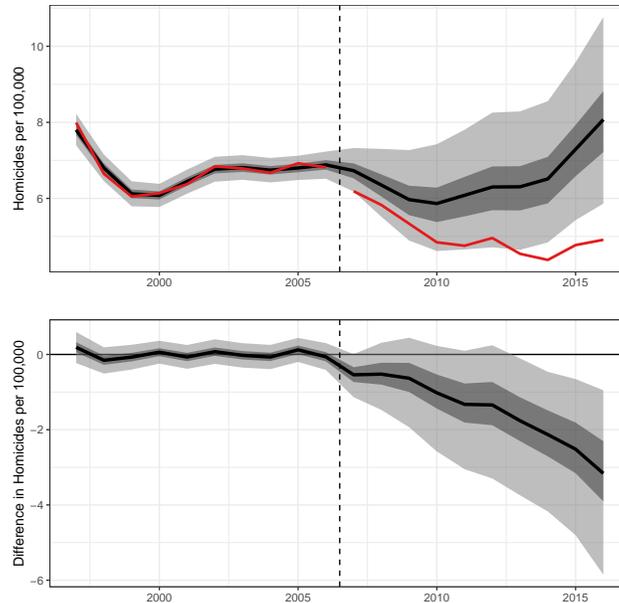


Fig 4: Impact of APPS on homicides per 100,000 in California by year. The observed data is show in red. The black line is the posterior mean estimate; the dark and light gray regions are, respectively, the 50% and 95% posterior predictive intervals.

the posterior mean as well as 50% and 95% posterior predictive intervals. The red line shows the observed murder rate for California. The estimates pre-treatment repeat the pre-treatment imbalance checks in Figure 3; the counterfactual estimates post-treatment are consistently higher than the observed murder rate. The lower pane shows this difference explicitly, where zero is no impact. These estimates suggest that APPS reduced murders in the state, with both impacts and uncertainty growing over time.

Figure 5 shows the posterior estimates for the average annual impact. The posterior mean estimate is roughly 1.5 murders avoided per 100,000 people per year, with a 95% credible interval of (0.32, 2.9) murders per 100,000 people; the posterior probability that the impact is negative is 99.6%. To contextualize these estimates, California’s annual murder rate in 2004 and 2005 was 6.8 per 100,000. Focusing on the smaller end of the 95% credible interval suggests a decline of at least 5 percent from baseline. We can also compare the annual number of murders prevented by this system to expenditures for APPS investigative teams in a single year, 2018, when California’s population was 39.6 million. A decrease in the murder rate of 0.32 corresponds to nearly 127 murders avoided per year, while the total funding for APPS investigation teams in fiscal year 2017-2018 was \$11.3 million (Petek, 2019). Dividing this cost of \$11.3 million by the estimated 127 avoided murders gives a (conservative) estimate of roughly \$89,000 per murder prevented. This is dramatically lower than the estimated benefit from the avoided crime, which is typically on the order of \$10 million (Heaton, 2010; Domínguez and Raphael, 2015). We include the posterior distribution of expenditures per average number of murders avoided in Figure 5.

One way to assess whether these estimates are potentially spurious is to conduct an “in-time” placebo check, where we set the APPS adoption year to be earlier than 2006 (Abadie, Diamond and Hainmueller, 2010). We can then compare the difference between the observed homicide rate and its posterior in the years between the placebo APPS year and 2006. Because we do not expect an effect in these held-out years, the difference should be near zero.

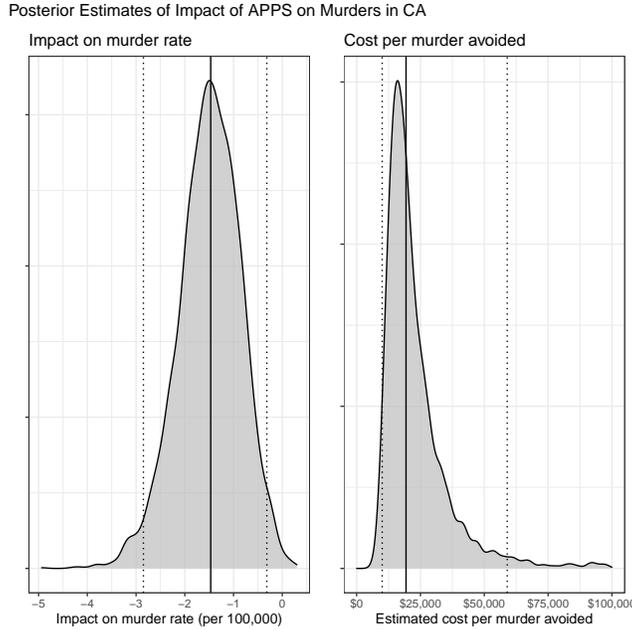


Fig 5: (left) Posterior distribution of the average impact of APPS on homicides per 100,000 in California for 2007 to 2016; (right) Posterior distribution of the cost per murder avoided, with a state population of 39.6 million and annual budget expenditure of \$11.3 million. The solid line denotes the posterior mean; the dashed lines denote 95% credible intervals.

Figure 6 shows these placebo estimates for placebo APPS over four years (2002-2005). We see that the posterior distribution is indeed near zero for all four placebo estimates, indicating that our preferred MTGP model finds no treatment effects when we expect there to be none. Note, however, that the relatively small number of pre-intervention years leads to several limitations to this placebo check. First, with the earlier placebo years there is a great deal of uncertainty due to the limited number of pre-intervention years. Second, we only evaluate placebo effects up to four years after the placebo APPS year. Settings with more time periods would have less uncertainty in the placebo check and would allow us to evaluate placebo effects for more years, which would enable us to assess the gradual increase in the magnitude of the treatment effect we discuss below.

Our preferred estimates for the impact of APPS are similar to the preliminary SCM estimates in Section 2.3. Both estimate that the counterfactual California would have seen the homicide rate drop more slowly and plateau at a higher level, before seeing a larger rise between 2014-2016. Both procedures estimate similar average impacts (a posterior mean of -1.5 vs. -1.41 for SCM), but the MTGP approach allows us to coherently quantify the uncertainty in the impacts, incorporating the fact that we are measuring impacts on homicide *counts*. Here we see that there is substantial uncertainty in the impact, but that there is a very high probability the impact is negative. In Appendix Figure C.12, we compare these estimates to those from various alternative estimators: (i) SCM without an intercept shift; (ii) Ridge Augmented SCM (Ben-Michael, Feller and Rothstein, 2021); (iii) fitting a low-rank factor model via `gsynth` (Xu, 2017); and (iv) low-rank matrix completion (Athey et al., 2021). These estimators broadly agree with our MTGP approach, with point estimates mostly within the 50% posterior predictive intervals and all well within the 95% posterior predictive intervals. Finally, we consider the relative impact of the other 49 comparison states on our estimates via a “leave one out” approach that removes each comparison state from the panel, finding

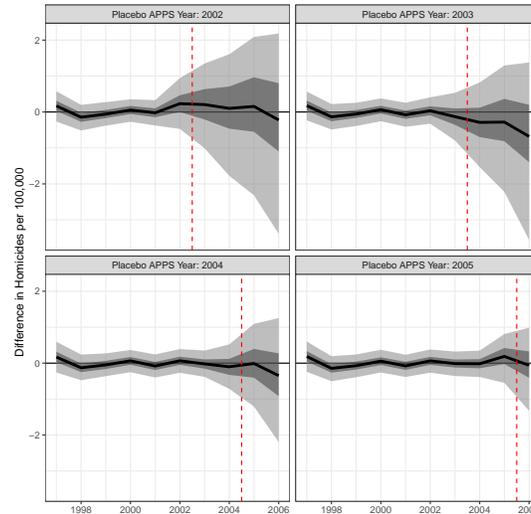


Fig 6: In-time placebo estimates for the rank 5 Poisson MTGP model, setting the placebo APPS year to be 2002-2005.

the posterior without conditioning on that state’s homicide rate series. We show the results in Appendix Figure C.13. Overall we find that removing each state has little effect on the posterior distribution; the largest differences come from removing Illinois, New Jersey, and Texas, but these estimates still broadly comport with our main estimates.

Our findings are also suggestive of larger effects in later years relative to earlier years. While these results may be due to an increase in murder rates in the comparison states, there are aspects of the APPS program that could lead to larger effects over time. First, there might be a lag in the effect of removing a firearm on violent incidents. Since the program reduces the stock of firearms over time, the effects may similarly increase along with the stock of firearms removed. Moreover, since people usually enter the APPS database upon purchasing a new gun, the proportion of guns and gun owners covered by the system is likely increasing over time as well.

California has also implemented further restrictions on firearms since 2006, including: banning the carrying of an unloaded handgun in public (2012); requiring new semi-automatic hand guns to have microstamping technology (2013); requiring tougher safety tests for first time buyers (2014); requiring background checks for ammunition and limiting ammunition sales to individuals with a registered firearm in the APPS system (passed in 2016 but implemented in June 2019); and most recently imposing a 21 year age limit for firearm purchases and a lifetime ban for prior convictions for a violent misdemeanor (Christopher, 2019). While some of these reforms occur after the end of our study period, it is clear the state has increased the regulation of firearms above and beyond what has occurred in other states. To the extent that some of these other policies have reduced murder rates, our estimates will reflect the joint effect of the APPS intervention along with these other policies. This also potentially contributes to the effect estimates increasing over time. That being said, many of these policy changes were only possible due to the underlying information and infrastructure that support the APPS enforcement effort.

6.2.2. *Separate impacts on gun- and non-gun-related homicides.* As we discuss in Section 2.3, one falsification check involves disaggregating homicides by whether a firearm was

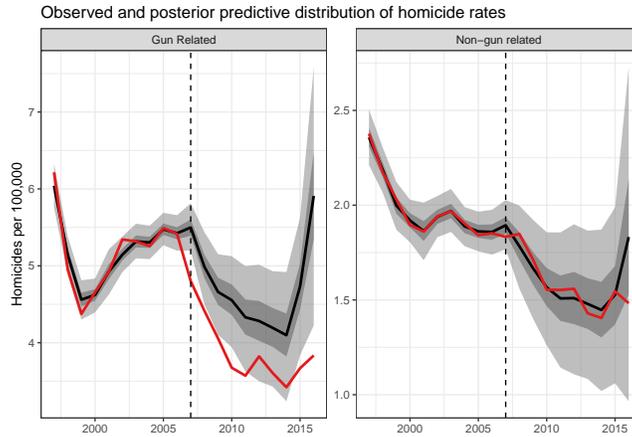


Fig 7: Impact of APPS on gun- and non-gun-related homicides per 100,000 in California by year. The observed data is show in red. The black line is the posterior mean estimate; the dark and light gray regions are, respectively, the 50% and 95% posterior predictive intervals.

used. Specifically, we do not expect that the APPS system would have reduced non-gun homicides. In fact, to the extent that there is substitution across weapons we might even expect that APPS may have increased homicide rates where firearms are not used.

To explore this, we use the [Kaplan \(2019\)](#) concatenation of the UCR Supplementary Homicide Files, restricted to our study period (1997 through 2016) and use the incident level data to calculate totals by state and year of the number of murders involving a firearm and the number of murders where a gun is not used.²¹ We then use these totals in conjunction with state level population counts to tabulate rates per 100,000.²² In general, gun murder rates are two to three times that of non-gun murder rates.

Figure 7 shows estimates for gun- and non-gun-related homicides, estimated via a joint multitask GP. Fitting both outcomes jointly is advantageous since these two outcomes are positively correlated due to shared latent factors. (See Appendix Figure C.14, for the posterior predictive correlation between gun and non-gun related homicides for California). While estimates are derived from a different data source than used above, the impacts for gun-related homicides are very similar to the estimated overall impacts in Figure 5. In contrast, the estimated impact of APPS on non-gun-related homicides is essentially zero.

7. Discussion. This paper uses a multitask Gaussian Processes framework in a panel data setting to estimate the causal effect of an innovative gun control program in California. MTGPs find a natural middle ground between pure “horizontal” time series forecasting models, based on extrapolating the treated unit’s outcomes series alone, and pure “vertical”

²¹We identify gun homicides as any incidents where the weapon used by any of the offenders in the incident are flagged as “firearm, type not stated,” “handgun,” “other gun,” “rifle,” or “shotgun.” Note, for a small percent of cases the weapon used is unknown (roughly 5 percent). We allot these observations to non-gun homicides, which should bias us towards finding an effect of APPS on non-gun homicides.

²²For these models, we restrict the donor pool to states where we observe homicide reports for all years from 1997 through 2016; inclusive of California, 43 states report data for all years. Murder victims in the SHR data by year and state often do not add to the officially reported titles in summary tables provided in the official UCR panel data. For some state-years, the ratio of the murder rate tabulated from SHR data to that tabulated from the UCR summary data is less than one while for others the value is above one. The median of this ratio is 1.02 and the mean is 1.005.

regression models, based on a weighted average of control units’ outcomes. Further, MT-GPs automatically quantify the relevant sources of uncertainty, providing a unified Bayesian framework in a setting where uncertainty quantification is notoriously challenging. They also allow for extensive model diagnostics via posterior predictive checks. Building on a robust GP literature, there are also many immediate extensions to the standard MTGP setup, including allowing for multiple outcomes and incorporating a mean model. Using this flexible approach, we find large effects of APPS on homicides in California, with program benefits far exceeding costs under typical assumptions.

There has been a flurry of policy activity since 2006 in the criminal justice domain, which could change crime rate patterns and the interpretation of our findings. First, since 2011 the state has cut back on using state prison for parole violations and for less serious felony offenses. Second, various state ballot initiatives have scaled back the severity of sentencing (Lofstrom, Bird and Martin, 2016). These reforms have reduced the overall incarcerated population (prisons and jails combined) of the state by roughly one-quarter, and may have led to higher crime rates through reverse incapacitation effects. However, several studies of these reforms fail to find any evidence of an impact of these changes on violent offenses (Lofstrom and Raphael, 2016; Bartos and Kubrin, 2018), and so we do not believe that our results are biased by these policy changes.

Methodologically, there are several promising directions for future research. First, there are many existing results for Gaussian Processes that would apply immediately in other settings but are not as relevant in our application. For instance, standard panel data models face a range of practical challenges when the data structure deviates from a so-called “balanced panel,” either due to irregularly sampled observations over time, missing data, or comparison states enacting policies similar enough to the policy of interest to warrant exclusion from the sample (see Imai and Kim, 2019). The MTGP framework naturally accommodates both irregularly sampled time points and missing values, and, unlike most Frequentist approaches, will automatically propagate the corresponding uncertainty.

Additionally, while California and the APPS program is uniquely comprehensive, several states have enacted forms of firearm surrender or removal policies for, e.g., those subject to a domestic violence restraining order. Therefore, an important extension of this approach is to allow for multiple treated units that adopt treatment over time, also known as staggered adoption (e.g., Ben-Michael, Feller and Rothstein, 2022). First, we can extend to the case where multiple treated units adopt treatment at the *same* time (i.e., simultaneous adoption) by changing the target to the *average* of the treated units. However, the extension to the case when units adopt over time (i.e., staggered adoption) is slightly more complicated. For instance, we could adapt the model to have a separate kernel for each treatment time, possibly including a hierarchical structure on the kernels (Flaxman et al., 2015).

Next, we can develop a framework for assessing sensitivity to departures from key modeling assumptions for the MTGP approach, especially the restriction that the time and unit kernels are separable. While there are many possible generalizations (see, for example Álvarez, 2017), these are inherently under-specified in settings like ours with relatively few units and time periods. Thus, a sensitivity analysis approach in the spirit of Franks, D’Amour and Feller (2019) is a natural avenue for appropriately quantifying uncertainty in this case.

Finally, our proposed MTGP approach implicitly puts the same weight on the error for California as on other states, even though California is the target unit. We can explore alternative approaches for parameterizing similarity across units and time, such as by imposing an “arrowhead” structure on the overall covariance, or by learning a richer dependence structure across units; see, for example, Li and Kontar (2020). Moreover, our approach focuses on modeling the control potential outcomes for California, in the spirit of finite sample causal inference — but allowing for potentially arbitrary treatment effects. We could also directly

model the treatment effect, including by incorporating a prior, as part of the MTGP framework.

Beyond the methodological contributions of the paper, the findings presented here suggest that California's approach to targeted firearm retrieval offers a potential policy option for other states to control gun violence. To be sure, the system is complex and undergirded by an elaborate information infrastructure and dedicated cadre of state workers including the database of gun owners and their firearms, the criminal justice databases that record triggering incidents, and the personnel needed to conduct the daily investigations and to retrieve firearms from prohibited persons. While few states currently have such an infrastructure in place, California's system draws upon databases of transactions involving licensed dealers that record gun sales throughout the country. Hence, other states could follow suit. Our results suggest that additional enforcement expenditures generate large benefits in terms of reduced murder rates.

We should note subsequent extensions of California's background checks system and the likely interaction with the APPS program. In 2019, California began to require point-of-sales background checks for the purchase of ammunition. Any current gun owners who are flagged by APPS as prohibited and still in possession of a firearm would be prevented from legally purchasing ammunition, as would any non-gun owner who is legally prohibited from owning a gun. As we discussed above, there are many gun owners in the state who are deemed prohibited but from whom state authorities are unable to retrieve all firearms. These additional restrictions likely further restrain the access of these individuals to lethal means. Future research should investigate the impact of this policy change.

A further development that may erode the effectiveness of APPS concerns the possibility that prohibited persons may be subverting the law by purchasing "ghost guns," illegal firearms often assembled from kits that one can legally purchase and that are not serialized and recorded by state authorities. California's Attorney General recently filed a motion for an injunction against companies that manufacture ghost-gun kits, in light of recent concerns regarding the use of these illegal firearms in crimes throughout the state ([California Department of Justice, 2022](#)). A potential adaptation of APPS to address this particular development might be to extend restrictions on ammunition sales to legal gun owners, with the further restriction that ammunition caliber at sale match the caliber of the firearms legally owned by the buyer.

Acknowledgments. We would like to thank Alyssa Bilinski, Phil Cook, Cass Crifasi, Alex D'Amour, Peng Ding, John Donohue, Naoki Egami, Max Gopelrud, Jess Kunke, Luke Miratrix, Jacob Montgomery, Jesse Rothstein, Yotam Shem-Tov, Elizabeth Stuart, Liyang Sun, and seminar participants at Polmeth 2021 and the University of Washington for helpful comments and discussion. We are especially grateful to Eli Sherman for his involvement in an earlier version of this project. A portion of this paper was previously circulated with the title "The Effect of a Targeted Effort to Remove Firearms from Prohibited Persons on State Murder Rates."

SUPPLEMENTARY MATERIAL

Supplement A: Additional results and implementation details

(doi: [COMPLETED BY THE TYPESETTER](#); .pdf). Additional results, comparisons with other methods, diagnostics, and other implementation details.

Supplement B: Replication code and data

(doi: [COMPLETED BY THE TYPESETTER](#); .zip). Complete replication code and data.

Funding. This research was supported in part by the Hellman Family Fund at UC Berkeley and by the Institute of Education Sciences, U.S. Department of Education, through Grant R305D200010. The opinions expressed are those of the authors and do not represent views of the Institute or the U.S. Department of Education.

REFERENCES

- ABADIE, A., DIAMOND, A. and HAINMUELLER, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program. *Journal of the American Statistical Association* **105** 493–505.
- ALAA, A. M. and VAN DER SCHAAR, M. (2017). Bayesian inference of individualized treatment effects using multi-task gaussian processes. In *Advances in Neural Information Processing Systems* 3424–3432.
- ÁLVAREZ, M. (2017). Multi-output Gaussian Processes. Gaussian Process Summer School 2017.
- ANEJA, A., DONOHUE III, J. J. and ZHANG, A. (2011). The impact of right-to-carry laws and the NRC report: lessons for the empirical evaluation of law and policy. *American Law and Economics Review* **13** 565–631.
- ANTONELLI, J. and BECK, B. (2021). Heterogeneous causal effects of neighborhood policing in New York City with staggered adoption of the policy.
- ARKHANGELSKY, D. and IMBENS, G. W. (2021). Double-robust identification for causal panel data models Technical Report, National Bureau of Economic Research.
- ARKHANGELSKY, D., ATHEY, S., HIRSHBERG, D. A., IMBENS, G. W. and WAGER, S. (2019). Synthetic difference in differences Technical Report, National Bureau of Economic Research.
- ATHEY, S., BAYATI, M., DOUDCHENKO, N., IMBENS, G. and KHOSRAVI, K. (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association* 1–15.
- BARRON, J. (2022). Could One County’s Success with ‘Red Flag’ Orders be a Model. *New York Times*.
- BARTOS, B. J. and KUBRIN, C. E. (2018). Can we downsize our prisons and jails without compromising public safety? Findings from California’s Prop 47. *Criminology & Public Policy* **17** 693–715.
- BEN-MICHAEL, E., FELLER, A. and ROTHSTEIN, J. (2021). The Augmented Synthetic Control Method. *Journal of the American Statistical Association* **116** 1789–1803.
- BEN-MICHAEL, E., FELLER, A. and ROTHSTEIN, J. (2022). Synthetic controls with staggered adoption. *Journal of the Royal Statistical Society. Series B: Statistical Methodology* **84** 351–381.
- BONILLA, E. V., CHAI, K. M. and WILLIAMS, C. (2008). Multi-Task Gaussian Process Prediction. In *Advances in Neural Information Processing systems* 153–160.
- BRANSON, Z., RISCHARD, M., BORNN, L. and MIRATRIX, L. W. (2019). A nonparametric Bayesian methodology for regression discontinuity designs. *Journal of Statistical Planning and Inference* **202** 14–30.
- BRODERSEN, K. H., GALLUSSER, F., KOEHLER, J., REMY, N. and SCOTT, S. L. (2015). Inferring causal impact using Bayesian structural time-series models. *Annals of Applied Statistics* **9** 247–274.
- CARLSON, D. (2020). Estimating a Counter-Factual with Uncertainty Through Gaussian Process Projection.
- CASTILLO-CARNIGLIA, A., KAGAWA, R. M., CERDÁ, M., CRIFASI, C. K., VERNICK, J. S., WEBSTER, D. W. and WINTEMUTE, G. J. (2019). California’s comprehensive background check and misdemeanor violence prohibition policies and firearm mortality. *Annals of epidemiology* **30** 50–56.
- CHEN, Y., GARNETT, R., MONTGOMERY, J. and PRATI, A. (2022). A Multi-Task Gaussian Process Model for Difference in Differences with Parallel(ish) Trends.
- CHENG, L.-F., DUMITRASCU, B., DARNELL, G., CHIVERS, C., DRAUGELIS, M., LI, K. and ENGELHARDT, B. E. (2020). Sparse multi-output Gaussian processes for online medical time series prediction. *BMC medical informatics and decision making* **20** 1–23.
- CHERNOZHUKOV, V., WÜTHRICH, K. and ZHU, Y. (2021). An exact and robust conformal inference method for counterfactual and synthetic controls. *Journal of the American Statistical Association* 1–16.
- CHRISTOPHER, B. (2019). How California Got Tough on Guns. Cal Matters, <https://calmatters.org/explainers/california-gun-laws-policy-explained/> accessed on October 11, 2019.
- COOK, P. J. and DONOHUE, J. J. (2017). Saving lives by regulating guns: evidence for policy. *Science* **358** 1259–1261.
- NATIONAL RESEARCH COUNCIL (2005). *Firearms and Violence: A Critical Review*. Washington, DC: National Academies Press.
- DING, P. and LI, F. (2018). Causal inference: A missing data perspective. *Statistical Science* **33** 214–237.
- DOMÍNGUEZ, P. and RAPHAEL, S. (2015). The role of the cost-of-crime literature in bridging the gap between social science research and policy making: Potentials and limitations. *Criminology & Public Policy* **14** 589–632.
- DONOHUE, J. J., ANEJA, A. and WEBER, K. D. (2019). Right-to-carry laws and violent crime: A comprehensive assessment using panel data and a state-level synthetic control analysis. *Journal of Empirical Legal Studies* **16** 198–247.

- DONOHUE, J. and BOULOUTA, T. (2019). The Assault Weapon Ban Saved Lives. *Stanford Law School Legal Aggregate*. October 15, 2019; <https://stanford.io/2MWNsrV>.
- DOUDCHENKO, N. and IMBENS, G. W. (2016). Balancing, regression, difference-in-differences and synthetic control methods: A synthesis Technical Report, National Bureau of Economic Research.
- FERMAN, B. and PINTO, C. (2021). Synthetic controls with imperfect pre-treatment fit. *Quantitative Economics*.
- FIEDLER, C., SCHERER, C. W. and TRIMPE, S. (2021). Practical and Rigorous Uncertainty Bounds for Gaussian Process Regression. In *Proceedings of the AAAI Conference on Artificial Intelligence* **35** 7439–7447.
- FLAXMAN, S., WILSON, A., NEILL, D., NICKISCH, H. and SMOLA, A. (2015). Fast Kronecker inference in Gaussian processes with non-Gaussian likelihoods. In *International Conference on Machine Learning* 607–616. PMLR.
- FRANKS, A., D’AMOUR, A. and FELLER, A. (2019). Flexible sensitivity analysis for observational studies without observable implications. *Journal of the American Statistical Association*.
- GELMAN, A., CARLIN, J. B., STERN, H. S., DUNSON, D. B., VEHTARI, A. and RUBIN, D. B. (2013). *Bayesian data analysis*. CRC press.
- GOBILLON, L. and MAGNAC, T. (2016). Regional policy evaluation: Interactive fixed effects and synthetic controls. *Review of Economics and Statistics* **98** 535–551.
- GOLDSTICK, J. E., CUNNINGHAM, R. M. and CARTER, P. M. (2022). Current causes of death in children and adolescents in the United States. *New England journal of medicine*.
- GOOVAERTS, P. et al. (1997). *Geostatistics for natural resources evaluation*. Oxford University Press on Demand.
- GREEN, E. L. and FERNANDEZ, M. (2018). Trump Wants to Arm Teachers. These Schools Already Do. *New York Times*.
- GUSTAFSON, P. (2010). Bayesian inference for partially identified models. *The International Journal of Biostatistics* **6**.
- HAZLETT, C. and XU, Y. (2018). Trajectory balancing: A general reweighting approach to causal inference with time-series cross-sectional data. Available at SSRN 3214231.
- HEATON, P. (2010). *Hidden in plain sight: What cost-of-crime research can tell us about investing in police* **279**. Rand Corporation.
- HEMENWAY, D. (2017). *Private Guns, Public Health*. University of Michigan Press.
- HENSMAN, J., MATTHEWS, A. and GHARAMANI, Z. (2015). Scalable variational Gaussian process classification.
- HUANG, B., CHEN, C. and LIU, J. (2019). GPMatch: A Bayesian Doubly Robust Approach to Causal Inference with Gaussian Process Covariance Function As a Matching Tool. *arXiv preprint arXiv:1901.10359*.
- HUANG, B., ZHANG, K. and SCHÖLKOPF, B. (2015). Identification of time-dependent causal model: A gaussian process treatment. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- IMAI, K. and KIM, I. S. (2019). On the Use of Two-way Fixed Effects Regression Models for Causal Inference with Panel Data.
- IMBENS, G. W. and RUBIN, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- JYLÄNKI, P., VANHATALO, J. and VEHTARI, A. (2011). Robust Gaussian process regression with a Student-t likelihood. *Journal of Machine Learning Research* **12** 3227–3257.
- KAGAWA, R. M., CASTILLO-CARNIGLIA, A., VERNICK, J. S., WEBSTER, D., CRIFASI, C., RUDOLPH, K. E., CERDÁ, M., SHEV, A. and WINTEMUTE, G. J. (2018). Repeal of comprehensive background check policies and firearm homicide and suicide. *Epidemiology* **29** 494–502.
- KANAGAWA, M., HENNIG, P., SEJDINOVIC, D. and SRIPERUMBUDUR, B. K. (2018). Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582*.
- KAPLAN, J. (2019). Jacob Kaplan’s Concatenated Files: Uniform Crime Reporting (UCR) Program Data: Supplementary Homicide Reports, 1976-2017. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2019-07-15. <https://doi.org/10.3886/E100699V7>.
- KARCH, J. D., BRANDMAIER, A. M. and VOELKLE, M. C. (2018). Gaussian process panel modeling-Kernel-based longitudinal modeling.
- KIM, S., LEE, C. and GUPTA, S. (2020). Bayesian synthetic control methods. *Journal of Marketing Research* **57** 831–852.
- LI, M. and KONTAR, R. (2020). On Negative Transfer and Structure of Latent Functions in Multi-output Gaussian Processes. *arXiv preprint arXiv:2004.02382*.
- LIU, L., WANG, Y. and XU, Y. (2020). A practical guide to counterfactual estimators for causal inference with time-series cross-sectional data. Available at SSRN 3555463.
- LOFSTROM, M., BIRD, M. and MARTIN, B. (2016). *California’s historic corrections reforms*. Public Policy Institute of California Sacramento.

- LOFSTROM, M. and RAPHAEL, S. (2016). Incarceration and crime: Evidence from California's public safety realignment reform. *The ANNALS of the American Academy of Political and Social Science* **664** 196–220.
- LOTT, J. R. JR and MUSTARD, D. B. (1997). Crime, deterrence, and right-to-carry concealed handguns. *The Journal of Legal Studies* **26** 1–68.
- MCCOURT, A. D., CRIFASI, C. K., STUART, E. A., VERNICK, J. S., KAGAWA, R. M., WINTEMUTE, G. J. and WEBSTER, D. W. (2020). Purchaser Licensing, Point-of-Sale Background Check Laws, and Firearm Homicide and Suicide in 4 US States, 1985–2017. *American journal of public health* **110** 1546–1552.
- MENCHETTI, F. and BOJINOV, I. (2020). *Estimating causal effects in the presence of partial interference using multivariate Bayesian structural time series models*.
- MIRATRIX, L. (2020). Using Simulation to Analyze Interrupted Time Series Designs. *arXiv preprint arXiv:2002.05746*.
- MODI, C. and SELJAK, U. (2019). Generative Learning of Counterfactual for Synthetic Control Applications in Econometrics. *arXiv preprint arXiv:1910.07178*.
- NAISH-GUZMAN, A. and HOLDEN, S. (2008). The generalized FITC approximation. In *Advances in neural information processing systems* 1057–1064.
- CALIFORNIA DEPARTMENT OF JUSTICE (2015). Legislative Report Number One – Calendar Year 2014: Armed and Prohibited Persons System. Sacramento, CA.
- CALIFORNIA DEPARTMENT OF JUSTICE (2016). Legislative Report Number Two – Calendar Year 2015: Armed and Prohibited Persons System. Sacramento, CA.
- CALIFORNIA DEPARTMENT OF JUSTICE (2017). Legislative Report Number Three – Calendar Year 2016: Armed and Prohibited Persons System. Sacramento, CA.
- CALIFORNIA DEPARTMENT OF JUSTICE (2018). Legislative Report Number One – Armed and Prohibited Persons System SB140 Legislative Report Number Four – Calendar Year 2017 (Revised 3/8/2018). Sacramento, CA.
- CALIFORNIA DEPARTMENT OF JUSTICE (2019). Legislative Report Number One – Armed and Prohibited Persons System SB140 Legislative Report Number Five – Calendar Year 2018. Sacramento, CA.
- CALIFORNIA DEPARTMENT OF JUSTICE (2022). Attorney General Bonta Files Motion for Preliminary Injunction Against Gun Retailers to Halt Sales of Illegal Ghost Gun Kits.
- OGANISIAN, A. and ROY, J. A. (2020). A Practical Introduction to Bayesian Estimation of Causal Effects: Parametric and Nonparametric Approaches. *arXiv preprint arXiv:2004.07375*.
- PANG, X., LIU, L. and XU, Y. (2020). A Bayesian Alternative to Synthetic Control for Comparative Case Studies. *Available at SSRN*.
- PEAR, V. A., WINTEMUTE, G. J., JEWELL, N. P. and AHERN, J. (2022). Firearm Violence Following the Implementation of California's Gun Violence Restraining Order Law. *JAMA network open* **5** e224216–e224216.
- PETEK, G. (2019). The 2019-2020 Budget: Analysis of the Governor's Criminal Justice Proposals. Sacramento, CA.
- PINKNEY, S. (2021). An Improved and Extended Bayesian Synthetic Control. *arXiv preprint arXiv:2103.16244*.
- RASMUSSEN, C. E. (2004). Gaussian processes in machine learning. In *Advanced lectures on machine learning* 63–71. Springer.
- RAY, K. and VAN DER VAART, A. (2018). Semiparametric Bayesian causal inference using Gaussian process priors. *arXiv preprint arXiv:1808.04246*.
- REN, B., WU, X., BRAUN, D., PILLAI, N. and DOMINICI, F. (2021). Bayesian Modeling for Exposure Response Curve via Gaussian Processes: Causal Effects of Exposure to Air Pollution on Health Outcomes. *arXiv preprint arXiv:2105.03454*.
- RICHARDSON, T. S., EVANS, R. J. and ROBINS, J. M. (2011). Transparent parameterizations of models for potential outcomes. *Bayesian Statistics* **9** 569–610.
- ROBINS, J. M., ROTNITZKY, A. and ZHAO, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association* **89** 846–866.
- RUBIN, D. B. (1973). Matching to remove bias in observational studies. *Biometrics* 159–183.
- RUBIN, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of statistics* 34–58.
- RUDOLPH, K. E., STUART, E. A., VERNICK, J. S. and WEBSTER, D. W. (2015). Association between Connecticut's permit-to-purchase handgun law and homicides. *American journal of public health* **105** e49–e54.
- SAMARTSIDIS, P., SEAMAN, S. R., PRESANIS, A. M., HICKMAN, M., DE ANGELIS, D. et al. (2019). Assessing the causal effect of binary interventions from observational panel data with few treated units. *Statistical Science* **34** 486–503.
- SAMARTSIDIS, P., SEAMAN, S. R., MONTAGNA, S., CHARLETT, A., HICKMAN, M. and ANGELIS, D. D. (2020). A Bayesian multivariate factor analysis model for evaluating an intervention by using observational time series data on multiple outcomes.

- SCHULAM, P. and SARIA, S. (2017). Reliable decision support using counterfactual models. In *Advances in Neural Information Processing Systems* 1697–1708.
- SHEPPARD, D. (1999). Promising Strategies To Reduce Gun Violence. US Department of Justice, Office of Juvenile Justice and Delinquency Prevention.
- SIEGEL, M., XUAN, Z., ROSS, C. S., GALEA, S., KALESAN, B., FLEEGLER, E. and GOSS, K. A. (2017). Easiness of legal access to concealed firearm permits and homicide rates in the United States. *American journal of public health* **107** 1923–1929.
- SOLIN, A. and SÄRKKÄ, S. (2020). Hilbert space methods for reduced-rank Gaussian process regression. *Statistics and Computing* **30** 419–446.
- SWANSON, J. W., NORKO, M. A., LIN, H.-J., ALANIS-HIRSCH, K., FRISMAN, L. K., BARANOSKI, M. V., EASTER, M. M., ROBERTSON, A. G., SWARTZ, M. S. and BONNIE, R. J. (2017). Implementation and effectiveness of Connecticut’s risk-based gun removal law: does it prevent suicides. *Law & Contemp. Probs.* **80** 179.
- STAN DEVELOPMENT TEAM (2021). Stan Modeling Language Users Guide and Reference Manual. Version 2.26.
- GIFFORDS LAW CENTER TO PREVENT GUN VIOLENCE (2019). <https://lawcenter.giffords.org/gun-laws/policy-areas/who-can-have-a-gun/categories-of-prohibited-people/>. Accessed on October 10, 2019.
- GIFFORDS LAW CENTER TO PREVENT GUN VIOLENCE (2022). Extreme risk protection orders.
- TUOMAALA, E. (2019). The Bayesian Synthetic Control: Improved Counterfactual Estimation in the Social Sciences through Probabilistic Modeling. *arXiv preprint arXiv:1910.06106*.
- WAINWRIGHT, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint* **48**. Cambridge University Press.
- WEBSTER, D. W. and WINTEMUTE, G. J. (2015). Effects of policies designed to keep firearms from high-risk individuals. *Annual review of public health* **36** 21–37.
- WILLIAMS, C. K. and RASMUSSEN, C. E. (2006). *Gaussian processes for machine learning* **2**. MIT press Cambridge, MA.
- WILSON, A. G., DANN, C. and NICKISCH, H. (2015). Thoughts on massively scalable Gaussian processes. *arXiv preprint arXiv:1511.01870*.
- WINTEMUTE, G. J., BECKETT, L., KASS, P. H., TANCREDI, D., STUDDERT, D., PIERCE, G., BRAGA, A. A., WRIGHT, M. A. and CERDÁ, M. (2017). Evaluation of California’s Armed and Prohibited Persons System: study protocol for a cluster-randomised trial. *Injury prevention* **23** 358–358.
- WINTEMUTE, G. J., PEAR, V. A., SCHLEIMER, J. P., PALLIN, R., SOHL, S., KRAVITZ-WIRTZ, N. and TOMSICH, E. A. (2019). Extreme risk protection orders intended to prevent mass shootings: a case series. *Annals of internal medicine* **171** 655–658.
- WITTY, S., TAKATSU, K., JENSEN, D. and MANSINGHKA, V. (2020). Causal Inference using Gaussian Processes with Structured Latent Confounders. In *Proceedings of the 37th International Conference on Machine Learning* (H. D. III and A. SINGH, eds.). *Proceedings of Machine Learning Research* **119** 10313–10323. PMLR.
- XU, Y. (2017). Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models. *Political Analysis* **25** 57–76.
- ZHU, X. (2005). *Semi-supervised learning with graphs*. Carnegie Mellon University.
- ZIMRING, F. E. (1972). The medium is the message: Firearm caliber as a determinant of death from assault. *The Journal of Legal Studies* **1** 97–123.