PNAS
www.pnas.org

Supplementary Information for

Changes in firearm mortality following the implementation of state laws regulating firearm access and use

Terry L Schell PhD*, Matthew Cefalu PhD, Beth Ann Griffin PhD, Rosanna Smart PhD, Andrew R. Morral PhD

Terry L Schell PhD
Email: tschell@rand.org

This PDF file includes:

Supplementary text
Figures S1 to S2
Tables S1 to S6
SI References

www.pnas.org/cgi/doi/10.1073/pnas.1921965117

Methods

State Laws. Our data on the implementation dates for the three classes of firearms legislation we analyzed were drawn from the RAND State Firearm Law Database *(1)*. Table S1 contains data on the years over which each law type applied to each state. For the purposes of the present paper, we code a state as having a child access prevention (CAP) law if the law specifies either civil or criminal penalties for storing a handgun in a manner that allowed access by a minor. We code a state as having a stand your ground (SYG) law if the state has a law that permits the use of lethal force for self-defense outside of the defender's home or vehicle, even when a retreat from danger would have been possible. Without such laws, individuals who use deadly force in self-defense may face criminal or civil penalties if they could have avoided the threat by leaving the situation or using non-deadly means for defense. We code a state as having a right to carry (RTC) law if concealed carry permits are issued whenever legally permissible without the discretion of law enforcement. Specifically, states that either prohibit concealed carry of firearms, or that "may issue" concealed carry permits are coded as not having a RTC law; states that either "shall issue" concealed carry permits to those who meet legal requirements or that allow concealed carry without any permit are coded as having a RTC law. Almost all the RTC laws being analyzed moved states from a regime in which concealed carry permits were issued to individuals only at the discretion of a law enforcement agency (May issue), into a regime in which all individuals who applied for a permit and met the legal requirements were issued a concealed carry permit (Shall issue).

Empirically, these three laws are associated. The adoption of RTC laws is positively associated with SYG laws and negatively associated with CAP laws. For example, while 36% of all states have adopted a CAP law, among the 18 states that have both RTC and SYG laws, only 18% have a CAP law.

Analytic Covariates. The model includes covariates that were the first 17 principle components extracted from a larger set of time-varying state characteristics. These characteristics have been found by other researchers to be associated with firearm deaths or are variables that are commonly used when analyzing state-level differences in health or crime. These variables are all taken from federal government sources and constitute descriptive statistics for each state for each year in the studied period. Additional information about each variable can be found in Schell, et al. *(2)*. The 34 characteristics are shown in Table S2.

The characteristics marked with * are plausibly the effects of gun control policies as well as possible confounds for estimating the policy's effects. For this reason, they are lagged one year in the model; they predict firearm deaths in the subsequent year, rather than in the year they are measured.

Prior to analysis, a few of these state characteristics were cleaned or transformed to mitigate undesirable distributional properties. Specifically, in a few cases in which values were missing for a given state-year, we imputed values using linear interpolation between the prior-year and subsequent-year values for that state. For a few predictors with extreme outliers, we also applied modest transformations to limit the influence of outlier values. Specifically, we applied the minimal power transformation (e.g., square root) that ensured all values were within five standard deviations of the mean.

Finally, we conducted additional transformations of the state characteristics to address the high degree of collinearity among some of these variables. We used dimension-reduction techniques to capture most of the information contained in the full set of correlated variables with a smaller number of orthogonal variables. To do this, we (a) removed from each covariate the variance that is collinear with the year effects that will be included in the models, and (b) used principal components analysis to linearly transform the matrix of standardized covariates into an ordered set of orthogonal variables (principal components). The models used the first 17 of these principal components, which were selected because those variables could explain more than 95 percent of the variance in the full matrix of state characteristics. Using this transformed matrix of covariates makes more efficient use of the available data; virtually all of the information contained in the original 35 state characteristics can be captured using only 17 model parameters. When included in the model, the 17 principal components were each standardized to mean = 0, standard deviation = 1.

Table S1. Implementation dates for CAP, SYG, and RTC laws by state.

State	CAP Law	SYG Law	RTC Law
Alabama	NA	Jun 2006 to present	NA†
Alaska	NA	Sep 2006 to present	Oct 1994 to present
Arizona	NA	Apr 2006 to present	Apr 1994 to present
Arkansas	NA	NA	Feb 1995 to present
California	Jan 1992 to present	NA	NA
Colorado	NA	NA	May 2003 to present
Connecticut	Oct 1990 to present	NA	NA
Delaware	Jul 1994 to present	NA	NA
Florida	Oct 1989 to present	Oct 2005 to present	Oct 1987 to present
Georgia	NA	Jul 2006 to present	Aug 1989 to present
Hawaii	Jun 1992 to present	NA	NA
Idaho	NA	NA	Apr 1990 to present
Illinois	Jan 2000 to present	NA	Jul 2013 to present
Indiana	NA	Jul 2006 to present	Jun 1983 to present
Iowa	Apr 1990 to present	NA	Jan 2011 to present
Kansas	NA	Jul 2011 to present	Jan 2007 to present
Kentucky	NA	Jul 2006 to present	Oct 1996 to present
Louisiana	NA	Aug 2006 to present	Apr 1996 to present
Maine	NA	NA	Sep 1985 to present
Maryland	Oct 1992 to present	NA	NA
Massachusetts	Oct 1998 to present	NA	NA
Michigan	NA	Oct 2006 to present	Jul 2001 to present
Minnesota	Aug 1993 to present	NA	Aug 2003 to present
Mississippi	NA	Jul 2006 to present	Jul 1990 to present
Missouri	NA	Aug 2007 to present	Feb 2004 to present
Montana	NA	Apr 2009 to present	May 1991 to present
Nebraska	NA	NA	Jan 2007 to present
Nevada	Oct 1991 to present	Oct 2011 to present	Jul 1995 to present
New Hampshire	Jan 2001 to present	Nov 2011 to present	NA†
New Jersey	Jan 1992 to present	NA	NA
New Mexico	NA	NA	Jan 2004 to present
New York	NA	NA	NA
North Carolina	Dec 1993 to present	Dec 2011 to present	Jul 1995 to present
North Dakota	NA	NA	Apr 1985 to present
Ohio	NA	NA	Apr 2004 to present
Oklahoma	NA	Nov 2006 to present	Jan 1996 to present
Oregon	NA	NA	Jan 1990 to present
Pennsylvania	NA	Aug 2011 to present	Jun 1989 to present
Rhode Island	Jun 1995 to present	NA	NA
South Carolina	NA	Jun 2006 to present	Aug 1996 to present
South Dakota	NA	Jul 2006 to present	Mar 1985 to present
Tennessee	NA	May 2007 to present	Oct 1996 to present
Texas	Sep 1995 to present	Sep 2007 to present	Jan 1996 to present
Utah	NA	Mar 1994 to present	May 1995 to present
Vermont	NA	NA	NA
Virginia	Jul 1992 to present	NA	Jul 1995 to present
Washington	NA	NA	NA†
West Virginia	NA	Feb 2008 to present	Aug 1989 to present
Wisconsin	Apr 1992 to present	NA	Jul 2011 to present
Wyoming	NA	NA	Oct 1994 to present

Note: †Alabama, New Hampshire, and Washington have right-to-carry laws with dates of enactment that substantially predate the time period we studied. Those laws are not considered in our analysis to ensure the estimated effects of laws are identified using within state changes.

Table S2. State characteristics included in analyses.

Type	Variable
Age Distribution (percentages)	
	Younger than 15
	15–29
	30–44
	45–60
	60–75
	Older than 75
Race or ethnicity (percentages)	
	Caucasian
	African American
	Asian/Pacific Islander
	American Indian/Alaska Native
	Hispanic
Relationship status (percentages)	
	Married/widowed
	Divorced
	Never married
Highest education (percentages)	
	Without high school diploma
	High school diploma
	Four-year degree
	Graduate degree
Other demographics	
	Total population size
	Gender ratio
	Percentage of children in single parent household
	Percentage foreign born
	Percentage military veterans
	Percentage urban households
	Percentage >25 years old, black, and urban
	Percentage >25 years old, Hispanic, and urban
Socioeconomic conditions	
	Percentage of population in the workforce
	Percentage unemployed (U3)
	Average income (inflation adjusted)
	Poverty rate
	Alcohol consumption per capita
	Incarceration persons per capita
	Police officers per capita
	Percent receiving hunting license*

Note: * because these variables were viewed as plausibly endogenous to state firearm regulations, these covariates were lagged by one-year.

Statistical Model and Law Coding. The selection of the model was based on a large-scale simulation designed to identify the most appropriate model specification for analyzing state-level firearm death data *(2)*, out of more than 100 methods investigated. The most appropriate model on the basis of the simulations demonstrated accurate type I error (i.e., unbiased standard errors), the best statistical power (i.e., lowest actually error of the estimate), and minimal bias (i.e., directional and magnitude bias was always a small fraction of the standard error of the estimate). The model used in the current research is a variant of the model selected through those simulations. Relative to that model, the treatment effects were modified to allow greater flexibility in phase-in function for any possible causal effect. This change reflects the fact that there is considerable uncertainty about the phase-in function of these laws, unlike in the simulation where it was known.

For simplicity, we present a version of the model below estimating the effect of a single law, although our final model included indicators for all three laws. Let y_{st} and N_{st} be the number of firearm deaths and the population size within a given state s in year *t*, respectively. Further, let $\boldsymbol{X}_{st}$ denote the set of 17 principal components from the covariates in state s and year *t*, and define an autoregressive variable $L_{s,t-1}$ equal to $\log\left(\frac{y_{s,t-1}}{N_{s,t-1}}\right)$. For a single law, this model can be expressed as:

$$y_{st} \sim NegBin(\mu_{st}, \phi)$$

$$\log(\mu_{st}) = \log(N_{st}) + \alpha + \zeta_t + \delta L_{s,t-1} + \boldsymbol{X}_{st}\boldsymbol{\gamma} + \beta_1 I_{st} + \beta_2\left(I_{st} - I_{s,t-1}\right) + \beta_3 A_{st} + \beta_4\left(A_{st} - A_{s,t-1}\right)$$

where ζ_t is a year-specific effect. The presence of a specific law is coded in I_{st} and A_{st}. I_{st} is coded to represent an instant effect of the law on the outcome. It is 0 for each state-year without the law, 1 for each state-year with the law fully in effect, and a pro-rated valued for years with partial implementation based on the number of months in which the law was in effect. A_{st} is coded to represent a gradual phase-in of the law's effect. It is computed as a spline that is linearly increasing from 0 to 1 over 72 months beginning with the month of law implementation. Because the outcome is modelled as discrete-time based on calendar year, yearly values of the spline used in the model are computed as the average of the monthly spline value over each calendar year.

The model includes these two codings of the law (instant and 6-year) as individual predictors, but they are also included in a first-differenced form, i.e., change in the level of the indicator variables since the prior year. When the autoregressive coefficient δ is zero, the model is a standard generalized linear model (GLM) predicting the level of the outcome; in that case the coefficients β_1 and β_3 are the appropriate unbiased estimators of the effect of the instant and phased-in effects of the law, respectively. When the autoregressive coefficient δ is one, the model becomes equivalent to a first-differences model, in which case β_2 and β_4 (the first-differences of the law indicators) are the appropriate estimators of the causal effect of the of the instant and phased-in effects of the law, respectively. The model includes both sets of indicators so that the model could fit the desired function relating the laws to the outcome for any value of the autoregressive coefficient between 0 and 1. For example, if the law has a true instant effect on the firearm death rate that persists indefinitely, the model can reproduce that function in each year regardless of how much of the effect in any given year is mediated through the autoregressive coefficient (because the full effect was already reflected in the prior year's value) versus being a direct effect of the law indicators. Essentially, the model we are using is a parent model from which both the standard levels models and first-differences (i.e., change) models fall out as special cases. Unlike both standard levels and first-differences models, however, the current model ensures that residuals are uncorrelated across adjacent years, which is a key assumption of the likelihood function for all these models.

Selection and implementation of priors. The priors used for the Bayesian estimation of the model are listed below. For the variables in the model other than the law indicators, all priors were intended to be either uninformative or weakly informative.

- $\alpha \sim N(-10, 100^2)$ -- the prior for the overall intercept is noninformative with a standard deviation of 100, covering all plausible base rates of the outcomes, and centered at a value of -10 to reflect the expected rarity of the outcomes.
- $\delta \sim Unif(0, 1)$ -- the weakly informative prior for the autoregressive coefficient is uniform and constrained to fall between 0 and 1.
- $\zeta_t \sim N(0, 0.65^2)$ -- there is a 95% prior probability that a given year's effect has an IRR between 0.28 and 3.55. This weakly informative prior was based on our assumption about the proportion of variance in the data that might be attributable to year effects conditioned on the other factors in the model.
- $\boldsymbol{\gamma} \sim N(0, 0.11^2)$ -- separately for each principal component of the covariates, there is a 95% prior probability that the IRR for a one standard deviation change in a given covariate is between 0.81 and 1.24. This weakly informative prior was based on our assumption about the proportion of variance in the data that might be attributable to these covariates conditioned on the other factors in the model.

The prior distributions on the laws' effects are more complex and reflect the interdependence among the β distributions and their dependence on the autoregressive coefficient δ. Conceptually, these priors are designed to ensure that the total marginal effect of each law, when integrated through the model, has a known prior distribution over time and across a range of values of δ.

- $\beta_1 \sim N\left(0, \frac{(1-\delta)\tau^2 b^2}{2}\left(\frac{1-\delta^{k+1}}{1-\delta}\right)^{-2}\right)$
- $\beta_2 \sim N\left(0, \frac{\delta\tau^2 b^2}{2}\right)$
- $\beta_3 \sim N\left(0, \frac{(1-\delta)\tau^2 b^2}{2}\left(\frac{1-\delta^{k+1}}{1-\delta} - \frac{1}{k+1}\sum_{t=1}^{k} t\,\delta^t\right)^{-2}\right)$
- $\beta_4 \sim N\left(0, \frac{\delta\tau^2 b^2}{2}\left(\frac{1-\delta^{k+1}}{(k+1)(1-\delta)}\right)^{-2}\right)$

Specifically, b^2 was chosen to ensure that the prior distributions imply the intended variance (more on the choice of b^2 below). We wanted to have a prior on the instant effect of the law, integrated over both the β_1 and β_2 effects and through the autoregressive term, to have the intended variance, b^2. To achieve that, we have priors on β_1 and β_2 that exchange variance such that when the autoregressive coefficient is zero the prior on β_1 has the desired variance while the prior on β_2 is set to zero. As the autoregressive coefficient approaches 1, this becomes a first-difference model and the prior on the β_2 reflects the desired variance while the prior on β_1 goes to zero. There is a comparable trade off in the variance across the priors for β_3 and β_4 designed to ensure the variance in the phased-in effect of the law, integrated through the model, has the desired variance, b^2.

These formulae for the prior distributions were derived by replacing the autoregressive term $L_{s,t-1}$ with its corresponding expected value (i.e., $\mu_{s,t-1}$). As a result, the formulae are only approximate, and an empirically derived scaling factor τ^2 was added to the approximation to ensure the proper variance of the marginal effect. The value of τ^2 was estimated through simulation and was set at $\tau^2 = 1.46^2$.

Because the process for setting these priors is complex and approximate, we simulated the total effect of a law under our prior distributions to verify it had the intended distribution. Specifically, we simulated the prior for firearm deaths (the total effect of the law in each year integrated through the model) for RTC laws. All parameters other than the βs were set to their posterior means. The empirically simulated prior distribution resulted in a year 6 effect with a median IRR=1.00 and 95% CI from 0.83 to 1.20. Note that the intervals are not symmetric around 1.00 because (a) the marginal effect is a complex, nonlinear function of the parameters and (b) we constructed credible intervals with equal probability in the two tails. Other credible intervals can be constructed that are symmetric.

That is to say, our prior reflects the belief there is a 50% chance that each policy increases deaths and a 50% chance it decreases deaths, but that it is unlikely that any of these policies increase firearms deaths by more than 20% or decrease them by more than 17%. This prior distribution was similar across years, with a constant median and an expected slight increase in dispersion over 6 years; the 95% CI for the year 1 IRR was 0.86 to 1.16.

This prior variance on the total effect of the law, which is set by the scaling parameter b^2, was selected to be consistent with expert opinion as assessed in a survey of gun policy experts *(3)*. In this survey, a diverse set of gun policy experts were asked about the anticipated effects of 13 policies on firearm homicides and suicides. All the studied policies in the survey were modest regulations that have been held to be constitutional and have been tried in states (including the effects of polices changing regulations for concealed carry and liability for improperly stored weapons, which are studied here). The prior distribution for the law effects used in this study were selected to reflect the aggregate opinion of the surveyed experts about the effect sizes of 13 similar firearm laws. Specifically, the target prior distributions for log(IRR)'s based on the expert surveys were N(0, 0.08), N(0, 0.07), and N(0, 0.09) for total firearm deaths, firearm suicide, and firearm homicide, respectively. Because we used an informative prior that implies the effects of these laws are unlikely to be descriptively large, we conducted a sensitivity test (included later in this document) in which we re-estimate effects with weakly informative priors.

Effect Estimation and Presentation. As discussed above, the effects of a given law on firearms deaths are represented across four separate coefficients in the model, and none of those coefficients can be interpreted individually to represent the unbiased effect of the law in any given year. To facilitate interpretation of this model, we compute a marginal effect of each law, integrated through the model, to yield a total effect of the law in each year following implementation.

Computing this marginal effect is more complicated than for many models for several reasons. First, autoregressive models are recursive such that a direct effect of the law on time T implies effects of the law at T+1 are mediated through the autoregressive path, even in the absence of a direct effect of the law on T+1. This mediated effect causes standard treatment indicators to yield biased effect estimates. Second, this is a nonlinear model, such that the marginal effect is a nonlinear transformation of the model parameters that depends on the levels of other covariates in the model. Third, the model assumes a negative binomial distribution, which, unlike a Normal distribution, cannot be ignored when computing the marginal effects in an autoregressive model. Because of these issues, it is possible that there are slight differences in the marginal effects computed in time periods with different values on the covariates. We chose to conduct the marginalization of the effect at the end of the study period under the assumption that the inferential population of interest is close to what would occur today or in the future based on a change in policy.

We use the following method to estimate the total marginal effect of the law for each year after implementation assuming that all states implemented the law in year 2008. Specifically, the effect of interest *d* years after implementation is the ratio of the expected total number of firearm homicides in year 2008 + *d* if all states implemented the law in year 2008 to the expected total number of firearm homicides in year 2008 + *d* if none of the states had the policy at any point in years 2008 through 2008 + *d*. These expectations are nonlinear functions of the model parameters and will be expressed through the following expectations:

$$\mu_{1,s,2008+} = E\left[y_{st} \mid I_{s,2008} = 1, \ldots, I_{s,2008+d} = 1\right]$$

and

$$\mu_{0,s,2008+d} = E\left[y_{st} \mid I_{s,2008} = 0, \ldots, I_{s,2008+d} = 0\right]$$

Note that N_{st} and X_{st} are assumed to be fixed and known for all *s* and *t*. The effect of interest, expressed as an IRR, is then given by:

$$\Delta_d = \frac{\sum_s \mu_{1,s,2008+}}{\sum_s \mu_{0,s,2008+d}}$$

Since there is no simple closed form for Δ_d as a function of the model parameters except at *d=0*, we estimate it using Monte Carlo simulation. For a given set of model parameters, do the following:

1. Set I_{st} and A_{st} based on the assumption that all states implemented the law in 2008.That is, set $I_{st} = 1$ for t from 2008 to 2016 and A_{st} based on the phase-in starting in 2008.
2. For *m=1 to M,*
   - Generate $(\tilde{y}_{s,2008}^{(m)} \mid L_{s,2007}) \sim NegBin(\mu_{s,2008}, \phi)$ with $\mu_{s,2008}$ as before.
   - For *d* = 1 to 8,
     a. Define $\tilde{L}_{s,2008+d-1}^{(m)} = \log\left(\frac{\tilde{y}_{s,2008+d-1}^{(m)}}{N_{s,2008+d-1}}\right)$
     b. Generate $(\tilde{y}_{s,2008+d}^{(m)} \mid \tilde{L}_{s,2008+d-1}^{(m)}) \sim NegBin(\mu_{s,2008+} \ , \phi)$ with $\mu_{s,2008+d}$ as before but using $\tilde{L}_{s,2008+d-1}^{(m)}$ as the autoregressive variable.
3. Take the sample average of $\tilde{y}_{s,2008+d}^{(m)}$ as an approximation of $\mu_{1,s,2008+}$
4. Repeat (1)-(3) for $\mu_{0,s,2008+d}$ by setting I_{st} and A_{st} based on the assumption that none of the states had the law from 2008 to 2016.
5. Calculate the effect of interest, $\Delta_d = \frac{\sum_s \mu_{1,s,2008+d}}{\sum_s \mu_{0,s,2008+d}}$

To get the posterior distribution of Δ_d, this process is repeated for each MCMC sample of the posterior distribution of the model parameters.

We also present the effect sizes as national count data, in addition to IRRs. This was done by getting an estimated number of deaths with and without a given law in each state in 2016. For states that had the law in 2016 we take their actual death count $y_{s,2016}$ as the known number of deaths with the law and use $y_{s,2016}/\Delta_d$ to get the posterior distribution of the number of deaths without the law. For states without the law in 2016 we take their actual death count $y_{s,2016}$ as the known number of deaths without the law and use $y_{s,2016} \times \Delta_d$ to get the posterior distribution of the number of deaths with the law. The national estimate for the change in the number of deaths is the sum of these state estimates with the law minus the sum of state estimates without the law.

When describing the posterior distribution of Δ_d we present the median, rather than the mean. This was done because it is not possible to have an unbiased mean on the prior distribution for both the model scale (log counts) and the original count scale. That is, if the log(IRR) of the prior has mean zero, the IRR cannot have mean 1. The posterior median, however, is unbiased in count and IRR units as well as in the model units. In addition to medians of the posterior distributions, we present the 80% credibility interval to describe the range in which the true effect is likely to be found, and we present the probability that the true effect is negative (i.e., that the law is beneficial) as the proportion of the posterior distribution less than IRR = 1.

Finally, the underlying data used to estimate the model is discrete time in which all data within a calendar year are combined (e.g., our outcome data and covariates are annual measures). Therefore, while the underlying spline functions in the model are continuous time functions, the model estimates are effectively discrete time estimates. Thus, when we label results in tables or figures as “in the 1st year after implementation,” the estimate should be interpreted as the average effect over the first 12 months following implementation. Our primary outcome was chosen to be the 6th year after implementation, which corresponds to the effect of the policy averaged over the 61st-72nd months following implementation, which differs slightly from the point estimate of the function at 72 months post implementation. We chose to use this time-point as our primary outcome to maximize a trade-off between problems with shorter and longer evaluation periods. Assessing the effect of the laws further from the implementation increases our confidence that we are assessing the long-term effect of the law, which is likely the primary concern for policy makers, rather than a temporary or transitory effect. However, estimating a causal effect many years after the hypothesized cause poses both statistical and conceptual challenges. The effect estimates themselves get more variable over time, particularly when some implementing states have not had the law long enough to contribute to the long-term effect estimate. Thus, those longer-term effects will, on average, have somewhat greater error. In addition, making causal attributions when the effect and the hypothesized cause are separated by lengthy intervals is more difficult as the range of potential confounding variables increases. The 6-year time period is well-within the observed period of these laws within the current data, which includes, on average, 23-years of data after implementation for CAP laws, 21-years for RTC laws, and 9 years for SYG laws.

Supplementary Results

Estimated Model. The models were estimated with 20,000 MCMC samples. The critical beta parameters all had more than 9300 effective samples, and all parameters showed Rhat values of approximately 1.00. Nominally, the model has 68 parameters. However, because of restrictions among the priors, as well as the use of weakly informative priors, the effective numbers of parameters were somewhat lower: 64, 65, and 63 for models of total firearms deaths, firearms suicides, and firearms homicides, respectively. The effective number of parameters is estimated using the calculation from the Wantanabe-Akaike Information Criterion *(4)*. The association between the model predicted rate of firearm deaths in a given state-year (posterior mean of predicted count ÷ population size) and the actual firearm death rate was $R^2 = 0.94$. Model predictions for the suicide and homicide subtypes of firearms deaths were similarly good fits to the data, $R^2 = 0.91$ and 0.94, respectively. This ability of the models to explain substantial variation in the data is largely due to the strong autocorrelation present in the data. The autoregressive coefficients were 0.90, 0.84. and 0.86 for total firearms deaths, firearms suicides, and firearms homicides, respectively, suggesting that the final model closely approximates a first-differences model with modest regression to the mean. Table S3 presents the posterior distribution for key parameters from the model of total firearm deaths.

Table S3. Posterior Distributions for Key Parameters in Model of Total Firearm Deaths.

Parameter	Posterior Median	95% CI lower bound	95% CI upper bound
alpha	-11.260	-11.310	-11.210
delta	0.896	0.876	0.917
Beta1_SYG	0.002	-0.008	0.011
Beta1_CAP	-0.001	-0.010	0.008
Beta1_RTC	0.001	-0.008	0.010
Beta2_SYG	0.032	-0.010	0.074
Beta2_CAP	-0.014	-0.056	0.028
Beta2_RTC	0.020	-0.011	0.051
Beta3_SYG	0.005	-0.008	0.018
Beta3_CAP	-0.004	-0.015	0.008
Beta3_RTC	0.002	-0.008	0.013
Beta4_SYG	-0.020	-0.117	0.075
Beta4_CAP	-0.058	-0.157	0.039
Beta4_RTC	0.008	-0.063	0.079

Note: Coefficients for year effects and covariates are omitted. Posterior medians for year effects varied between -0.085 and 0.038; Posterior medians for coefficients on standardized covariates varied between -0.025 and 0.011.

Sensitivity of findings to priors. To address any concerns that the conclusions of the paper depend on poorly chosen priors for the effects of laws, we replicated the primary results presented in Table 1 using weakly informative priors (Table S4). Specifically, we estimated the models with priors on each of the four betas of $N(0, 0.4^2)$. Practically, this is more than 25 times the prior variance on those effects relative to our preferred model specification and reflects a prior that these laws might have very large effects on firearms deaths (given the actual autoregressive coefficient, an IRR of 5 is well within the prior). The result of this sensitivity test shows that all effect sizes became slightly more extreme with less informative priors, with the modal change in

IRRs = 0.01 across the two sets of priors. The changes for the homicide outcome were slightly larger, which is anticipated because firearm homicide was the sparsest outcome with the widest credibility intervals. Similarly, the posterior probabilities that the policies reduced deaths became somewhat more extreme with less informative priors. For example, the posterior probability that the restrictive policy regime was associated with a reduction in firearm homicide increased from 0.95 with our preferred priors to 0.97 with less informative priors. Our preferred priors thus result in probability estimates that have somewhat greater uncertainty in the direction of the effects.

Table S4. Effects of CAP, RTC, and SYG laws on change in state firearm death rates in the 6th year after implementation, estimated using minimally informative priors

Outcome Law	Posterior median IRR	80% CI lower bound IRR	80% CI upper bound IRR	Posterior probability of reduced deaths
Firearm Deaths				
CAP	0.92	0.88	0.97	0.98
RTC	1.04	1.00	1.07	0.11
SYG	1.03	0.98	1.08	0.21
Restrictive	0.87	0.79	0.94	0.99
Firearm Suicides				
CAP	0.94	0.90	0.99	0.93
RTC	1.04	1.00	1.08	0.08
SYG	1.02	0.97	1.07	0.34
Restrictive	0.89	0.82	0.97	0.96
Firearm Homicides				
CAP	0.90	0.82	0.98	0.95
RTC	1.03	0.97	1.10	0.24
SYG	1.07	0.99	1.17	0.15
Restrictive	0.81	0.70	0.93	0.97

Phase in of the joint effective of restrictive laws. The main report describes the effect size at the 6th year and the general shape of the phase in when a restrictive regime of laws is implemented. This restrictive regime consists of CAP laws but no SYG or RTC law. Figure S1 illustrates the phase in of this regime's effects.

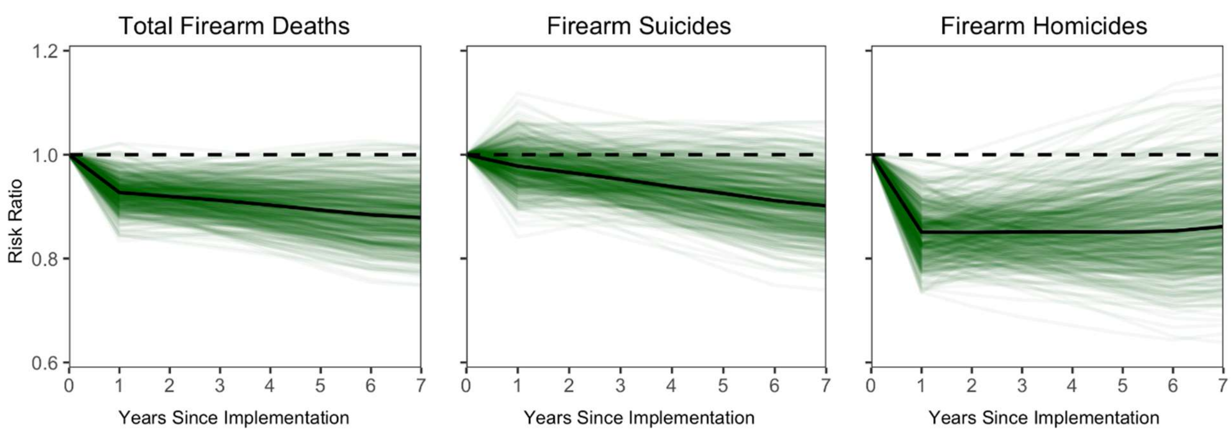

Fig S1. Posterior distribution of the effect of a restrictive legal regime (CAP laws, but no RTC, or SYG laws) over time, by type of firearms deaths. Effects are expressed as IRRs. Posterior median and 500 samples from the posterior distribution are plotted.

Effects of laws on total homicides and suicides. One concern about focusing exclusively on firearm death rates is that it may miss unintended or second-order consequences of firearm legislation. For example, some individuals prevented from committing a homicide or suicide using a firearm may commit the homicide or suicide using some other means. In that case, looking exclusively at firearms deaths may overstate the effect of the laws. On the other hand, homicide and suicide may show some form of social contagion in which these events make subsequent homicides or suicides of other victims more likely. In that case, looking exclusively at firearms deaths may underestimate the total effect of firearms laws. To address these limitations in looking exclusively at firearm deaths, we conducted a sensitivity test in which we ran our primary model on both total suicide and total homicide (including both firearm and non-firearm mechanisms). These estimates are presented in Table S5, and generally show a modest attenuation of the effects found on firearm homicide and suicide, which is what would be expected if these laws have minimal effects on non-firearm homicides and suicides.

A more direct assessment of these effect can be seen when comparing the predicted effects of the laws expressed as counts of homicides and suicides nationally (Table S6). These effects on firearms deaths and all deaths from homicide and suicide are generally quite close to each other and show a similar pattern across laws. In some instances, such as the CAP effect on suicide, the total estimated reduction in deaths is somewhat greater than the effect on firearm deaths alone. In other cases, such as the CAP effect on homicide, the total estimated reduction in deaths is slightly smaller than the effect on firearm deaths alone.

Table S5. Effects of CAP, RTC, and SYG laws on change in state suicide and homicide rates in the 6th year after implementation

Outcome Law	Posterior median IRR	80% CI lower bound IRR	80% CI upper bound IRR	Posterior probability of reduced deaths
Total Suicides				
CAP	0.96	0.93	0.99	0.95
RTC	1.02	0.99	1.04	0.17
SYG	1.00	0.97	1.03	0.53
Restrictive	0.94	0.89	0.99	0.92
Total Homicides				
CAP	0.96	0.91	1.02	0.82
RTC	1.01	0.97	1.06	0.35
SYG	1.04	0.98	1.10	0.20
Restrictive	0.91	0.83	1.00	0.89

Table S6. Mean change in the number of deaths nationally (with 80% credible interval) for CAP, RTC, and SYG laws by mechanism, in the 6th year after implementation

Outcome Law	Firearm Deaths	Total Deaths
Suicides		
CAP	-1075 (-2163, -17)	-1880 (-3339, -431)
RTC	771 (-16, 1530)	826 (-279, 1901)
SYG	343 (-663, 1314)	-61 (-1456, 1293)
Homicides		
CAP	-1060 (-2113, -56)	-768 (-1916, 337)
RTC	470 (-347, 1245)	271 (-630, 1121)
SYG	718 (-300, 1682)	736 (-377, 1811)

Effects of CAP laws on firearm suicides by minors. For our primary analyses we assess the relationship of CAP laws to total firearm deaths, as well as firearm homicide and suicide in the full population. Some other research on CAP laws look specifically at deaths of children as an outcome. We chose to examine the full population for two reasons. First, CAP laws generally mandate safe gun storage for all households that contain minors, and thus the law could affect the likelihood of firearm deaths for anyone in such a household and anyone who might be a victim of homicide by a gun from such a household. Conceptually, the effects of CAP laws are not restricted to deaths involving minors, even if that was the legislative intent of the law. Secondly, opponents of CAP laws may view these laws as harmful because they make firearms less effective as deterrents to crime and diminish the potentially protective effects of the weapons for their owners. Looking at the law's effect exclusively among children would not capture these hypothesized harmful effects of the law, and thus would ignore the concerns of those who oppose these laws.

However, the law is targeted to households with children and some might expect that its causal effects on deaths in which a child used a gun would be larger than the effect found in the general population. Unfortunately, it is not possible to get reliable data on the number of firearm homicides in which gun user was a child, however we can subset data on firearm suicides by the age of the victim. The estimated effect of CAP laws on firearm suicides of individuals 0-19 years old was very similar to that found on the full population. In the 6th year following implementation, the posterior median IRR for CAP was 0.93 (80% CI: 0.89 to 0.97), a slightly larger reduction than estimated in the full population (IRR=0.95). The phase in of the effect was slightly different among minors relative to the full population, with the effect on victims 0-19 years old staying nearly constant between IRR = 0.94 and 0.93 through seven years. In contrast, the effect size estimated within the general population increased steadily from IRR = 0.98 to 0.95 over the same period. Although the estimated effect size among young people is nominally larger, particularly in the first few years of implementation, given the variance in these posteriors these differences should be interpreted cautiously.

Timing of the shifts in death rates within implementing states. The autoregressive model used to estimate the effects of these laws is a dynamic model of change in deaths; it identifies the effects of laws primarily through the pattern of year-over-year changes in a state's death rates in the years immediately following implementation, and it allows for a parametric but flexible phase-in of the law's effect (See Figure 1). However, to insure the model was not detecting shifts in death rates that occurred prior to the laws, we investigated if there were meaningful patterns in the outcome series during the pre-implementation period for the implementing states that were not accounted for by the model. Specifically, in Figure S2 we plot the model residuals for the implementing states relative to the date of implementation, covering the 7 years before and after the implementation year. Like our effect estimates, these residuals are expressed as risk ratios, i.e., $\sum Y / \sum \hat{Y}$ summed over the implementing states for a given law and year. For reference, the figure also presents the estimated treatment effect implied by the model over the 7 years after implementation on the same scale. Inspection of these figures revealed no evidence of meaningful trends or reliable deviations from RR=1 in the pre-law periods.

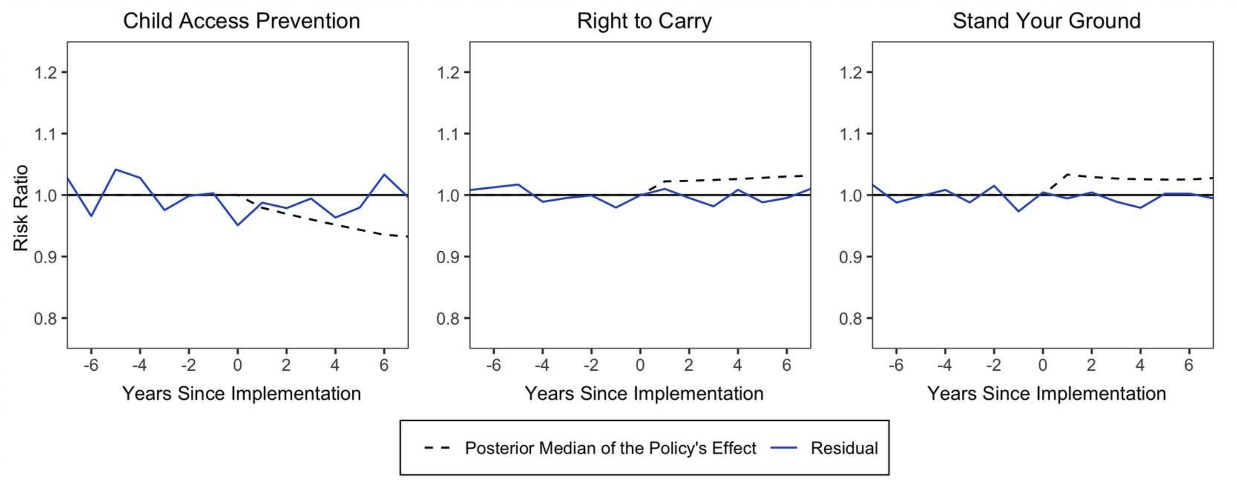

Fig S2. Model residuals for implementing states by time-relative-to-implementation for each law. Residuals are expressed as risk ratios ($\sum Y/\sum \hat{Y}$) and the estimated policy effects from Figure 1 are presented for reference on the same scale.

SI References

1. S. Cherney, A. R. Morral, T. L. Schell, *RAND State Firearm Law Database*. (RAND Corporation, Santa Monica, CA, 2018).
2. T. L. Schell, B. A. Griffin, A. R. Morral, *Evaluating Methods to Estimate the Effect of State Laws on Firearm Deaths*. (RAND Corporation, Santa Monica, CA, 2018).
3. A. R. Morral, T. L. Schell, M. Tankard, *The Magnitude and Sources of Disagreement Among Gun Policy Experts*. (Rand Corporation, Santa Monica, CA, 2018).
4. A. Gelman, J. Hwang, A. Vehtari. Understanding predictive information criteria for Bayesian models. *Statistics and Computing* **24**, 997-1016 (2014).