Patterns

Article

From Regression Analysis to Deep Learning: Development of Improved Proxy Measures of State-Level Household Gun Ownership

Graphical Abstract

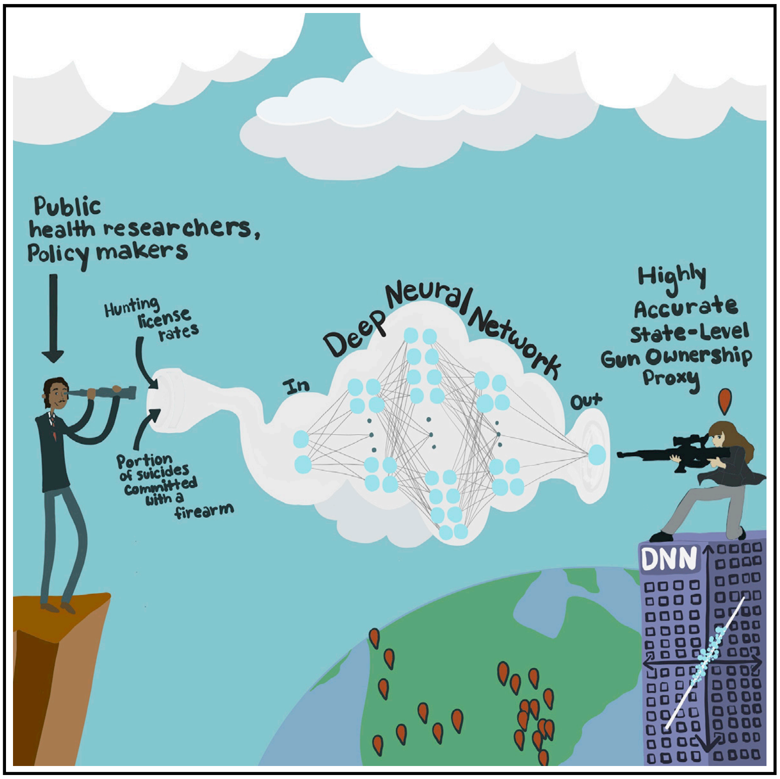

Highlights

- We develop deep-learning and regression-based proxies of state-level gun ownership
- Both new proxies significantly outperform existing ones
- We found that the widely used FS/S proxy is highly biased and inadequate
- We recommend FS/S no longer be used to represent state-level gun ownership

Authors

David Benjamin Gomez, Zhaoyi Xu, Joseph Homer Saleh

Correspondence

dbgomez94@gmail.com

In Brief

In the absence of direct measurements of state-level household gun ownership, the quality and accuracy of proxy measures for this variable are essential for firearm-related research and policy development. In this work, we develop two significantly improved proxy measures of state-level household gun ownership via two methods: first using regression analysis, then using deep learning. We subject the new regression-based and deep-learning proxies to critical examination, and we benchmark our new proxies against existing ones for accuracy and validity.

Gomez et al., 2020, Patterns *1*, 100154
December 11, 2020
https://doi.org/10.1016/j.patter.2020.100154

CellPress

Patterns

CellPress
OPEN ACCESS

Article

From Regression Analysis to Deep Learning: Development of Improved Proxy Measures of State-Level Household Gun Ownership

David Benjamin Gomez,[1,2,*] Zhaoyi Xu,[1] and Joseph Homer Saleh[1]
[1]School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA
[2]Lead Contact
*Correspondence: dbgomez94@gmail.com
https://doi.org/10.1016/j.patter.2020.100154

THE BIGGER PICTURE Data on state-level household gun ownership is largely missing in the United States, yet this variable is essential for firearm-related research and policy development. In the absence of gun ownership data, researchers and policy-makers have had to rely on proxy measures to represent this indispensable variable. Historically, the portion of suicides committed with a firearm has been regarded as the best proxy measure of gun ownership. In this work, we challenge this notion and develop two significantly improved proxy measures using first, traditional regression analysis and then the tools of deep learning. Our new proxy measures are both highly accurate and easy to obtain, and they can be used for a variety of purposes in cross-sectional studies of firearm-related violence at the state level.

1 2 3 4 **5** Mainstream: Data science output is well understood and (nearly) universally adopted

SUMMARY

In the absence of direct measurements of state-level household gun ownership (GO), the quality and accuracy of proxy measures for this variable are essential for firearm-related research and policy development. In this work, we develop two highly accurate proxy measures of GO using traditional regression analysis and deep learning, the former accounting for non-linearities in the covariates (portion of suicides committed with a firearm [FS/S] and hunting license rates) and their statistical interactions. We subject the proxies to extensive model diagnostics and validation. Both our regression-based and deep-learning proxy measures provide highly accurate models of GO with training R^2 of 96% and 98%, respectively, along with other desirable qualities—stark improvements over the prevalent FS/S proxy (R^2 = 0.68). Model diagnostics reveal this widely used FS/S proxy is highly biased and inadequate; we recommend that it no longer be used to represent state-level household gun ownership in firearm-related studies.

INTRODUCTION

The National Research Council (NRC) identified in its report on Firearm-Related Violence[1] a set of high-priority research topics to help inform the "development of sound policies that support both the rights and responsibilities central to gun ownership in the United States." The report noted that high-quality data are essential for the advancement of research and policy development, and without which "it is virtually impossible to answer fundamental questions" related to firearm-related violence. There is a complex and significant public health burden of firearm-related morbidity and mortality in the United States. This includes homicides, suicides—which receive much less public scrutiny but far outnumber homicides—and unintentional injuries and fatalities. Yet essential data on measures of exposure in the population are unavailable. For example, the report noted that "basic information about gun possession [and] ownership is lacking," and that there is a pressing need to develop accurate information about this topic. This observation echoes findings from a previous NRC report,[2] which noted that "the absence of information about gun ownership has been a major stumbling block for […] studies of suicide as well as for studies of homicide and other gun-related crime." In the absence of such information, and with Federal restrictions on certain data collection related to gun ownership,[3] researchers will have to continue to rely on proxy measures for this indispensable variable. As a result, the quality and accuracy of these proxy measures are essential for firearm-related research and policy development. Our work fits within this context and contributes toward this objective of developing more accurate proxy measures of gun ownership for public health research and policy development.

More specifically, gun ownership has been a main predictor or control variable in several important studies of suicide and gun violence.[4–21] When available, surveys conducted in the United States provide the most straightforward measure of household gun ownership—that is, the percentage of households with at least one gun. There is a handful of national surveys conducted in the United States that provide estimates of household gun ownership, for example, the General Social Survey (GSS) and the Gallup Poll surveys, but only one stratifies the surveys by state: the Behavioral Risk Factor Surveillance System (BRFSS).[22] Both the size and quality of the BRFSS has prompted experts in this field to refer to it as the proverbial gold standard. For example, in 2002 the BRFSS had roughly 230,000 observations of household gun ownership across 49 states, whereas the GSS had a nationally representative sample of only 2,800 respondents. More recently, the Gallup Poll surveys from 2008 to 2017 conducted their surveys with nationally representative samples between 300 and 5,000 respondents. Despite the BRFSS being the gold standard for state-level household gun ownership, it is not without its limitations. The decision to include the gun ownership question is made at the state level, and from 1995 to 2018 only one year (2001) included the gun ownership question on the surveys in all 50 states. The following year in 2002, 49 states included the question. For most years, few to no states included the gun ownership question. The surveys conducted in 2001 and 2002 constitute the most comprehensive two consecutive years in which the gun ownership question was included in the survey, and it is the standard by which the proxies developed in this work are trained and, as we will see, cross-validated.

In past studies, a myriad of proxy measures has been proposed for state-level household gun ownership. Two important works by Azrael et al.[23] and Kleck[24] reviewed these proxies and assessed their validity. Both concluded that the portion of suicides committed with a firearm (FS/S) is the best proxy measure of gun ownership for cross-sectional studies across large geographic regions. This proxy has been used extensively in studies of firearm-related violence.[25–36] The NRC report in 2005[2] noted that "many researchers have accepted FS/S as the best and, in fact, a nearly ideal proxy for studying the cross-sectional relationship between firearms and violence."

In this work, we propose that this measure be retired and no longer used as a proxy measure of state-level household gun ownership. It is not appropriate for the purpose and manner in which it is being used—as an independent variable for inference and prediction purposes, not subject to uncertainty or random variation—when it is in fact markedly biased and exhibits high variability, as we will show next. In its stead, we develop and evaluate two new proxies for state-level household gun ownership and discuss two different methods for obtaining them. The first is incremental and builds on a recently developed proxy that combines FS/S with state-level hunting license rates (HLR)[37] (and later used elsewhere).[38,39] We leverage traditional regression analysis to develop an improved proxy measure that accounts for non-linearities in these variables and their statistical interaction. There is no reason, theoretical or otherwise, to assume that gun ownership is best modeled as a linear combination of FS/S and HLR as opposed to a combination of transformations of these predictors and their statistical interaction. This possibility is explored in our first model-building strategy. We subject our regression model to critical examination and common diagnostics, and we benchmark our new proxy measure against the existing ones for accuracy and validity. Next, we leverage advances in deep learning (DL) to develop and validate a deep neural network (DNN) model of state-level household gun ownership based on the same two inputs, FS/S and HLR. We subject the DNN model to critical examination and benchmark our DL-based proxy against the regression-based proxies. The DL-based proxy significantly outperforms the others and is both highly accurate and largely bias free, as we will see next.

Our objectives in this work are 2-fold: (1) to provide the research community with accurate, easily obtainable proxy measures of state-level household gun ownership; and (2) to encourage broader adoption of advanced machine learning (ML) methods that go beyond the prevailing linear models in public health research. Significant opportunities reside in the adoption of ML in general, and DL in particular, for more accurate modeling of a host of public health problems, for teasing out novel insights from datasets, and ultimately for better-informed decision making.

RESULTS

Our results are based on a training dataset composed of a single response variable: state-level estimates of household gun ownership from the BRFSS surveys, averaged over years 2001–2002; and two covariates: (1) FS/S and (2) state-level HLR. Regarding the older proxies, we refer to the ubiquitous FS/S proxy as *Model 1* and the linear FS/S with HLR proxy as *Model 2*. We refer to our new regression-based proxy measure as *Model 3* (by design, *Model 3* uses the same covariates as those in *Model 2*), and we refer to the DL-based proxy as the *DNN*. For a fair comparison, all models were trained with the same 2001–2002 dataset and validated with BRFSS data from other years.

The New Regression-Based Proxy Measure

The first method by which we developed a new proxy measure of state-level household gun ownership was with traditional regression analysis. Our new proxy measure, *Model 3*, is incremental and builds on the older *Model 2* by accounting for the possibility of non-linearities and statistical interaction of the covariates FS/S and HLR. We examined different powers of FS/S and HLR and their statistical interaction to determine which combination provided the highest R^2 while maintaining statistical significance of the coefficients ($p < 0.05$). The details of our model development process are described in Method 1: Regression Analysis.

The results of our exploration of non-linear transformations of the covariates in *Model 3* are shown in Figure 1. Note that the

CellPress
OPEN ACCESS

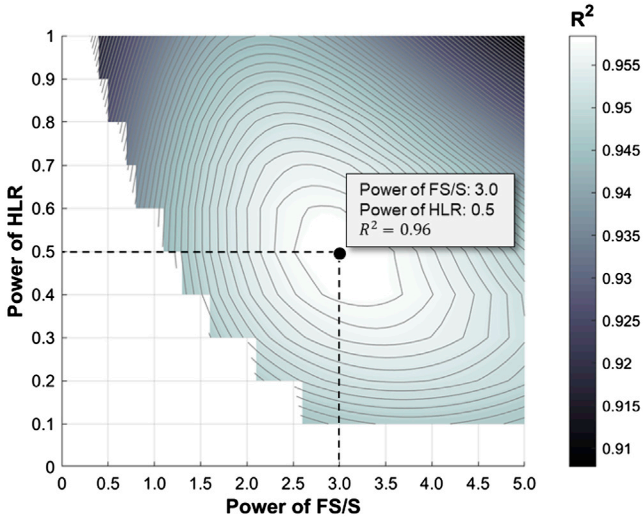

Figure 1. Results of the Exploration of Non-linear Transformations of the Covariates to Determine the Optimal Combination of Powers of FS/S and HLR in *Model 3*

interaction term is taken to be the product of the transformed covariates. The white space corresponds to combinations of coefficients that are not all statistically significant.

Figure 1 shows that the best powers of FS/S and HLR are 3.0 and 0.5, respectively. Our *Model 3* is given by the following expression:

$$\widehat{\mathrm{GO}} = \begin{array}{l} 1.046(\mathrm{FS/S})^3 + 0.094(\mathrm{HLR})^{0.5} - \\ 0.160(\mathrm{HLR})^{0.5}(\mathrm{FS/S})^3 - 0.011 \end{array}. \quad \text{(Equation 1)}$$

The New DL-Based Proxy Measure

The second method by which we devised a new proxy measure of state-level household gun ownership was with the tools of DL via a DNN. DL is a set of powerful methods with multiple layers of transformation, starting with the input variables and proceeding to more abstract levels of representation. "With the composition of enough such transformations, very complex functions can be learned," and intricate patterns or structures can be discovered in large datasets that are beyond the capabilities of regression or non-hierarchical models.[40] We developed a DNN with five layers: an input layer with two inputs, three hidden layers (each with a corresponding Leaky ReLU activation layer),[41] and an output layer with a single output. The hidden layers consist of 500, 1,000, and 500 neurons, respectively. Our DNN thus employs a 2-500-1,000-500-1 neuron distribution over each of the five layers, respectively; it takes in vectors of FS/S and HLR and outputs the predicted household gun ownership. We used the dropout method and L_2 regularization to avoid overfitting,[42] and we adopted the mean squared error (MSE) as the cost function with an ADAM[43] optimizer in the backpropagation process. Finally, we used the *k*-fold cross-validation technique to avoid overfitting and to assess how well the performance of the DNN will generalize to independent datasets (not used in the training of the DNN).[44,45] The details of our model development process are described in Method 2: Deep Learning.

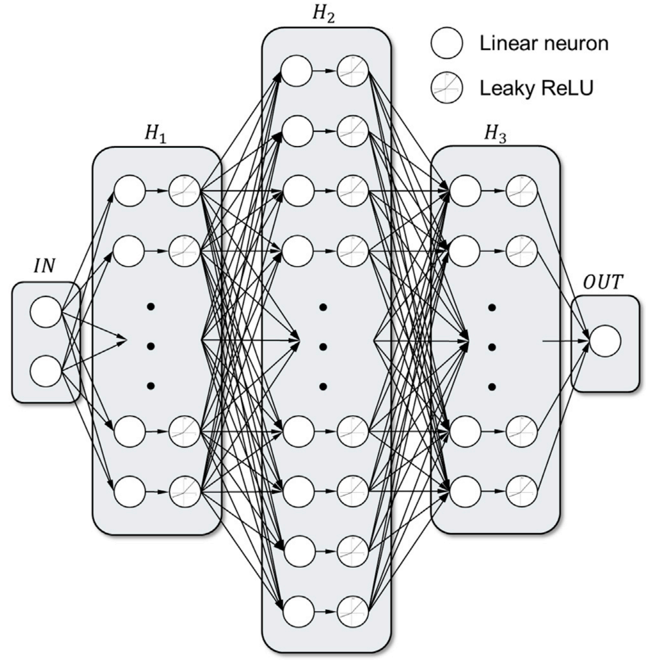

Figure 2. The DNN Architecture for the Development of an Improved Proxy Measure of State-Level Household Gun Ownership

H1, H2, and H3 refer to the first, second, and third hidden layer, respectively.

An important tradeoff arises in the design of the DNN architecture and the selection of its neuron distribution. On the one hand, increasing the number of neurons will improve the accuracy of the DNN for the training dataset. On the other hand, a DNN with too many neurons would likely suffer from overfitting, in which case the predictive capabilities on new datasets (generalization) would degrade. A good neuron distribution is one that achieves accurate performance on both the training and validation sets. We selected the DNN structure that simultaneously (1) minimized the performance gap on the training and validation sets and (2) achieved high accuracy on both sets (to avoid overfitting and underfitting). This was the 2-500-1,000-500-1 neuron distribution noted previously with three hidden layers. Our final DNN model structure is depicted in Figure 2. The DNN is freely available to be downloaded as a Python package and used (see GitHub). The repository includes the code and a ReadMe file that instructs a new user through the simple process of downloading and using the DNN either to reproduce our results or to run the model with a new dataset.

Comparative Analysis between the Old and New Proxy Measures of State-Level Household Gun Ownership

In this subsection, we conduct a comparative analysis of the old and new proxy measures of state-level household gun ownership in terms of model performance, diagnostics, validation, and predictive accuracy.

Model Performance

A visual summary of the training performances of the old and new proxy measures of state-level household gun ownership is shown in Figure 3. The regression model coefficients and

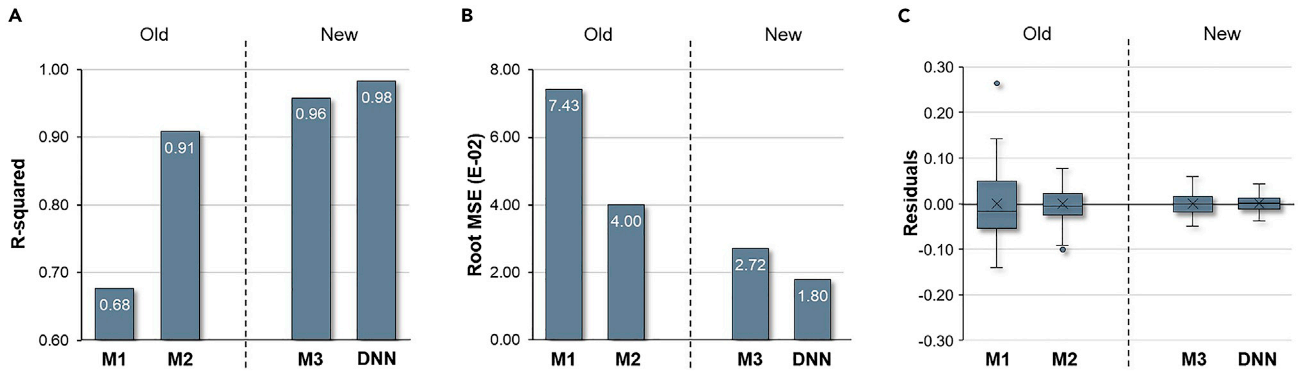

Figure 3. Performance Comparison of the Old and New Proxy Measures of State-Level Household Gun Ownership

From left to right, the figure depicts (A) the R^2, (B) the root MSE, and (C) the boxplot of the residuals ($e_i \equiv GO_i - \widehat{GO}_i$). M1, M2, and M3 refer to *Models 1*, *2*, and *3*, respectively.

corresponding statistics for all regression-based proxies are provided in Table 1. Recall that the FS/S proxy is referred to as *Model 1* and the linear FS/S with HLR proxy as *Model 2*. Our new regression-based proxy, which accounts for non-linearities in the covariates and their statistical interaction, is *Model 3*, and the new DL-based proxy is the DNN.

Whether measured by the R^2, the root MSE (RMSE), or the span and lack of symmetry of the residuals, the results in Figure 3 demonstrate that both new proxy measures, *Model 3* and the DNN, significantly outperform the older proxies, *Models 1* and *2*. Consider, for example, the residuals of the ubiquitous FS/S proxy, *Model 1* (Figure 3C). There is significant bias, with positive skew, in this model along with multiple outliers, including a prominent one at 27% where the model severely underestimates gun ownership (see the left-hand side of Figure 3C). *Model 2*, which incorporates HLR in addition to FS/S, represents a clear performance improvement over *Model 1*. Note first that 91% of the variability in the gun ownership data is explained by these two predictors in *Model 2*, in stark contrast with the limited "expressivity" of *Model 1*, which captures only 68% of said variability. However, *Model 2* remains biased, albeit to a lesser extent than *Model 1*, with inaccuracies up to 10 percentage points in estimated gun ownership.

In contrast, both the new regression-based proxy (*Model 3*) and the DNN provide significantly more accurate models of state-level household gun ownership, with R^2 of 96% and 98%, respectively. Their residuals are more symmetrically distributed (unbiased) and have a narrower span, as seen in Figure 3C.

Model Diagnostics

For transparency in model assessment, we conducted model diagnostics in the form of normal-quantile plots of the residuals, plots of residuals-versus-predicted gun ownership, and Cook's distances for influence analysis. These results are provided in Figures 4, 5, and 6, respectively. Note that Figures 4 and 5 contain roughly the same information, but we include both for visual clarity.

The normal-quantile plots in Figure 4 and the residual-versus-prediction plots in Figure 5 demonstrate the reduction in non-normality of the residuals between each model. In both Figures 4 and 5, we see a clear violation of the normality assumption (and homoscedasticity) with *Model 1*, whereas the DNN provides the most accurate and unbiased proxy measure with quasi-normally distributed residuals. Furthermore, consider the Cook's distances for influence analysis in Figure 6. (As a reminder, the Cook's distance diagnoses the robustness of a regression model to perturbations in the covariates.)[46] These results demonstrate that the parameters in *Model 1* are disproportionately influenced by a single observation or leverage point (South Dakota). Similarly, the parameters in *Model 2* are heavily determined by three observations (Montana [MT], Wyoming [WY], and South Dakota [SD]). In other words, fluctuations in the covariate data (FS/S and HLR) for those states during the years over which the models were trained (2001–2002) would have a disproportionate effect on the model parameters. In contrast, the Cook's distances for *Model 3* reveal that no one observation (state) has undue influence on the model parameters, i.e., the contributions from each state in the model training process are more evenly distributed, and the model parameters are less susceptible to fluctuations in the covariates. The results in Figure 6 indicate that the regression model with statistical interaction (*Model 3*) is more robust than the model without this term (*Model 2*). We further comment on the robustness and accuracy of *Model 1* after we discuss model validation and confidence intervals.

Model Validation

We conducted validation of all models using the gun ownership data from BRFSS surveys over other years for which data were available. We first validated on the survey results from 2004, as this represents the most complete independent dataset with ~280,000 observations over 49 states (see Table 3). We then combined the limited survey data from 1995 to 2017 (14 independent states in total) into our second validation set, although we acknowledge that the results from this set have to be taken cautiously given the small sample available (which is likely non-random). The R^2 and RMSE for the training set and both validation sets are provided in Table 2. As a side note, in the *k*-fold cross-validation of our DNN, we obtained an R^2 of 0.94 on folds not used in the training of the model (the cross-validation sets).

These results indicate that the DNN model remains robustly accurate across both validations and outperforms all the other models by a fair margin. Furthermore, a stark underperformance of the ubiquitous FS/S proxy (*Model 1*) is clear in Table 2. In the

Table 1. Model Coefficients and Associated Statistics for Both the Old (*Models 1* and *2*) and New (*Model 3*) Regression-Based Proxy Measures of State-Level Household Gun Ownership

Gun Ownership	Coefficient (95% CI)	SE	t	p > |t|
Model 1				
FS/S	0.847 (0.677 to 1.017)	0.0844	10.03	<0.001
Constant	−0.099 (−0.195 to −0.003)	0.0478	−2.07	0.044
Model 2				
FS/S	0.600 (0.498 to 0.703)	0.0508	11.82	<0.001
HLR	0.008 (0.006 to 0.009)	0.0007	10.88	<0.001
Constant	−0.047 (−0.100 to 0.006)	0.0262	−1.80	0.078
Model 3				
$(FS/S)^{3.0}$	1.046 (0.869 to 1.223)	0.0878	11.91	<0.001
$(HLR)^{0.5}$	0.094 (0.082 to 0.105)	0.0057	16.54	<0.001
Interaction	−0.160 (−0.209 to −0.110)	0.0246	−6.49	<0.001
Constant	−0.011 (−0.043 to 0.021)	0.0159	−0.68	0.500

FS/S, portion of suicides committed with a firearm; HLR, hunting license rates; CI, confidence interval; SE, standard error.

more complete validation set in 2004, the FS/S proxy explained a mere 61% of the total variability in gun ownership, whereas the DNN explained 94% of said variability. We also observe how well the *k*-fold cross-validation technique predicted the DNN's performance on independent datasets.

Confidence and Prediction Intervals

We examined more closely the predictive accuracy of all our regression-based proxies by computing the 95% confidence and prediction intervals (CI and PI, respectively). These results depict the extent to which the proxy measures would under- and overpredict state-level household gun ownership. We highlight first the FS/S proxy (*Model 1*); the regression model, along with the 95% CI and 95% PI, are shown in Figure 7.

Consider the 95% prediction interval. This is the interval of the estimated gun ownership for a single state with a given FS/S. We observe a roughly 15-percentage-point difference between the estimated gun ownership and the upper or lower bounds of the 95% prediction interval (flaring of these intervals notwithstanding). For example, for a state with FS/S = 0.5, *Model 1* predicts a gun ownership of roughly 32%; however, according to the 95% PI, the actual level of gun ownership for that state can be as high as 48% or as low as 17%. Thus, researchers who invoke FS/S as a proxy measure for state-level household gun ownership (e.g., to control for gun ownership) would be well advised to conduct uncertainty quantification and demonstrate that their findings are robust to states having this much uncertainty in their predicted levels of household gun ownership. Furthermore, the uncertainty in the FS/S proxy is not merely theoretical; we see from Figure 7 that there are several states with FS/S equal to roughly 0.5 for which the actual level of gun ownership (as provided by the BRFSS, 2001–2002 average) varies between 20% and 43%. To further illustrate this point, we compared the predicted and actual levels of household gun ownership for each state and for all regression-based proxies, as shown in Figure 8. Note that the data presented in Figure 8 are ordered according to the predicted levels of gun ownership and are plotted with the 95% PI.

Consider *Model 1* in Figure 8. There are several instances where there is a significant discrepancy between the actual levels of household gun ownership and those predicted by the model. For example, South Dakota (SD) and Maryland (MD) share similar predicted levels of household gun ownership, 30% and 31%, respectively, yet the actual levels of gun ownership for those states (as provided by the BRFSS) are 57% and 21%, respectively. Other instances of this weakness with *Model 1* can be found between Iowa/California (IA/CA) and North Dakota/Oregon (ND/OR). *Model 2* reduces the prediction intervals down to roughly 10 percentage points, but it remains biased and is disproportionately influenced by Montana (MT), South Dakota (SD), and Wyoming (WY). Among the regression-based proxies, the new *Model 3* provides the most accurate estimates of state-level household gun ownership with a prediction interval down to roughly 6 percentage points (again, flaring notwithstanding). The DNN model further halves this prediction interval.

DISCUSSION

In the absence of direct measurements of state-level household gun ownership, the quality and accuracy of proxy measures for this variable are essential for firearm-related research and policy development. In this work, we provide the research community with two new proxy measures of state-level household gun ownership. These new proxies are easily obtainable and provide significant improvements over the existing ones in terms of accuracy, reduced bias, and correlation with the variable they represent. Model diagnostics reveal that the widely used FS/S proxy measure is inadequate, and we recommend that it no longer be used to represent state-level household gun ownership in studies of firearm-related violence.

Our first new proxy builds on another recently developed proxy,[37] which combines FS/S with state-level HLR. We explored for, and identified, non-linearities in these variables and their statistical interaction. This increased the R^2 from 91% to 96%, reduced the model bias, and removed the high leverage points in the previous model, which overall improved the accuracy and robustness of the new model. We devised this new regression-based proxy for three reasons: (1) to provide public health researchers with an easily calculable proxy

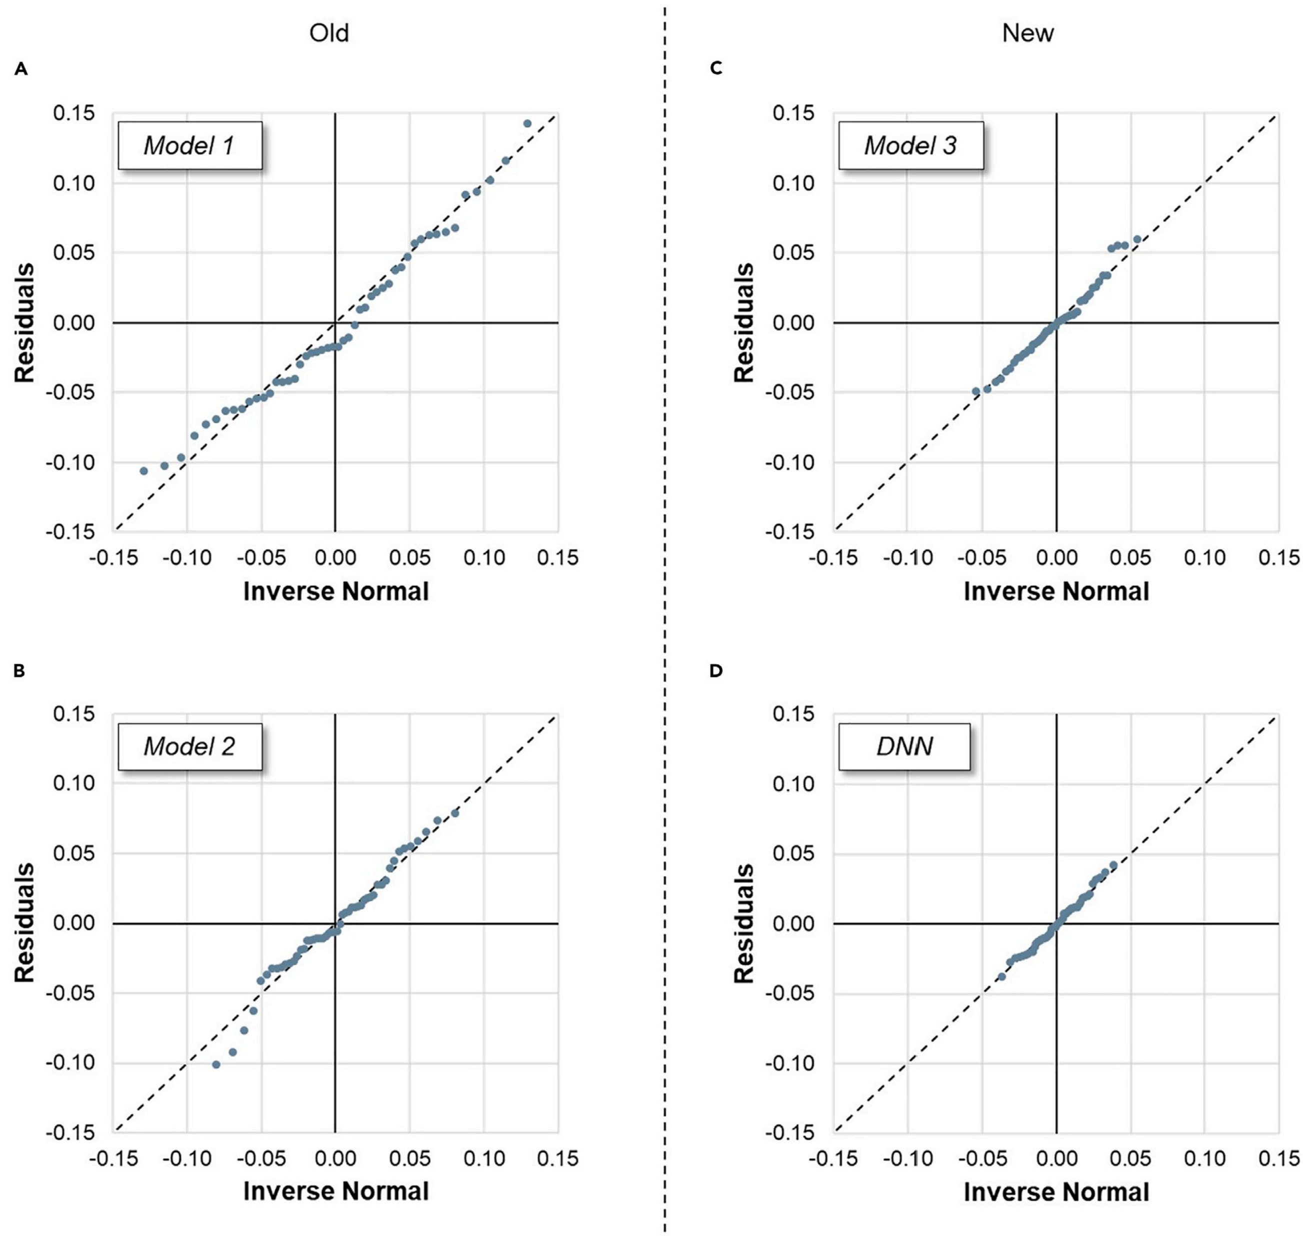

Figure 4. Normal-Quantile Plots of the Residuals for Both the Old and New Proxy Measures of State-Level Household Gun Ownership
Normal-quantile plots for (A–C) *Models 1–3* and (D) the DNN.

measure of state-level household gun ownership with Equation 1 using basic spreadsheets or simple computational tools; (2) to illustrate a useful distinction in statistical modeling between predictor and regressor (a single predictor x_i can have multiple regressors $\psi(x_i)$ in a model, $x_i \rightarrow \psi(x_i)$, e.g., polynomial terms of order p in x_i, and multiple predictors can be combined into a single regressor, $x_i,\ x_j \rightarrow \psi(x_i,\ x_j)$, e.g., an interaction term). This distinction is helpful to keep in mind when seeking to build better statistical models, and seems to be overlooked in some of the firearm-related literature. Finally, we did this to (3) emphasize the importance of performing model diagnostics for the examination of possible model deficiencies.[46] There is a dearth of model diagnostics in the literature on firearm-related violence. We believe this community would be well served if it adopted some statistical diagnostic tools, which would not only reveal novel insights and point the way toward better model building but also help meet expectations of transparency for model performance assessment.

For our second new proxy, we leveraged advances in DL to develop and validate a DNN model of state-level household gun ownership with the same two inputs used in the regression-based proxy. Our DNN proxy is the most accurate (R^2 = 0.98), unbiased

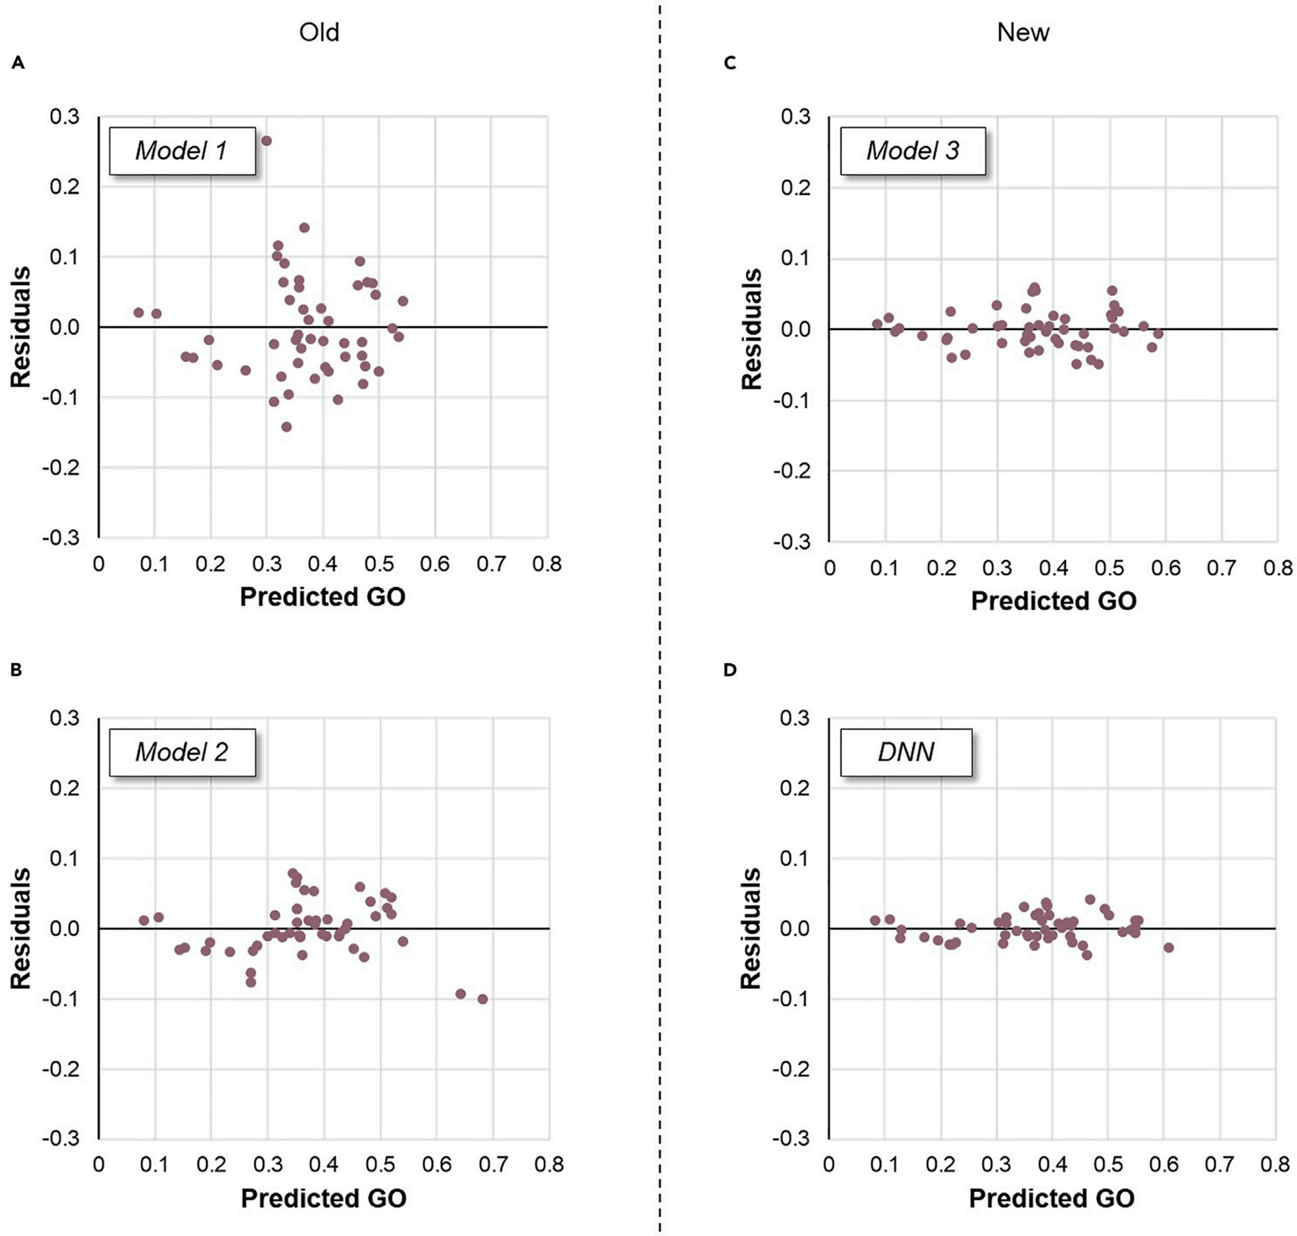

Figure 5. Plots of Residuals-versus-Predicted Gun Ownership for Both the Old and New Proxy Measures of State-Level Household Gun Ownership

Residuals-versus-predicted gun ownership (GO) plots for (A–C) *Models 1–3* and (D) the DNN.

model of state-level household gun ownership, and we recommend that it be used over any other proxy. The DNN model is freely available to be downloaded as a Python software package and used (see GitHub). The DNN model raises the following interesting challenge and requires a slightly different way of thinking from traditional statistical modeling: in regression analysis it is uncomplicated to see analytically and interpret how the inputs are transformed into the response variable; in contrast, in DL, with thousands (if not millions) of hierarchical transformations of the input vector(s), it is impractical to track these transformations.[47] With this level of complexity, one is confined to dealing with the model as a black box. Given the remarkable accuracy and robustness that can be achieved by DNN models, sacrificing some model transparency for significant performance improvement is a worthwhile tradeoff to make and adjust to. Furthermore, we demonstrated how the *k*-fold cross-validation technique can be used to avoid overfitting and assess how well the model will perform on a dataset not used in the training of its parameters.

We validated both of our new proxies using the 2004 BRFSS survey data on state-level household gun ownership as well as

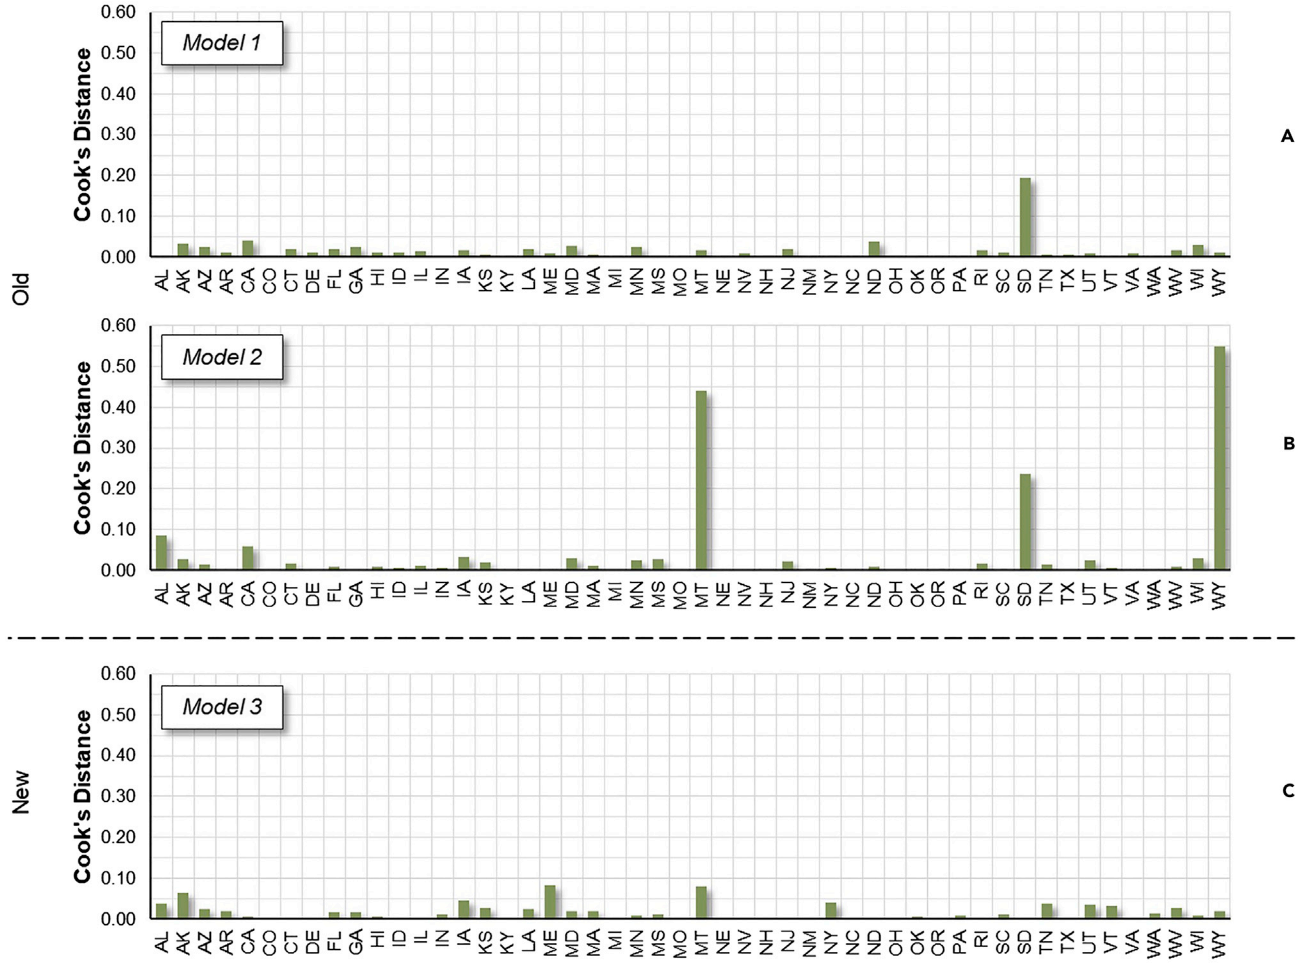

Figure 6. Cook's Distances for the Old and New Regression-Based Proxy Measures of State-Level Household Gun Ownership
Cook's Distances for (A) *Model 1*, (B) *Model 2*, and (C) *Model 3*.

with the limited data from 1995/2017. In both validations, we found that both of our new proxy measures outperformed the older ones, with the DNN providing the most accurate (and consistent) proxy measure for household gun ownership. We further examined the predictive accuracy of all regression-based proxies. We found that the FS/S proxy exhibits significant prediction uncertainty, with 95% PI roughly three times that of our new regression-based proxy. We emphasize that researchers who utilize this (or any) proxy measure of state-level household gun ownership would be well advised to examine and demonstrate whether their findings are robust to the uncertainty inherent in their proxy's predicted levels of household gun ownership.

Our results should be considered in light of several limitations they share with all other similar modeling studies. First, we assumed that the measures of gun ownership from the BRFSS are accurate representations of the actual levels of household firearm prevalence. This may not be the case given the sampling

Table 2. Model Validation Comparison of Both the Old and New Proxy Measures of State-Level Household Gun Ownership

		Training (50 States)		Validation #1 (49 States)		Validation #2 (14 States)	
		R^2	RMSE ($\times 10^{-2}$)	R^2	RMSE ($\times 10^{-2}$)	R^2	RMSE ($\times 10^{-2}$)
Old	*Model 1*	0.68	7.43	0.61	8.36	0.82	5.93
	Model 2	0.91	4.00	0.90	4.08	0.90	4.54
New	*Model 3*	0.96	2.72	0.92	3.62	0.90	4.53
	DNN	0.98	1.70	0.94	3.58	0.95	2.97

RMSE, root-mean-squared error.

Table 3. Available BRFSS Survey Data with Observations on State-Level Household Gun Ownership

Year	1995	1996	1997	1998	2001	2002	2004	2017
No. of states	11	7	7	4	50	49	49	3
No. of observations	21,754	15,530	14,744	9,526	201,409	231,180	284,384	17,825

These figures omit US Territories and the District of Columbia, as well as non-responses.

variability and uncertainties due to non-responses and erroneous responses inherent in the BRFSS. Second, the BRFSS survey results in 2001 and 2002 are nearly two decades old, and the structure of the relationship between gun ownership and both input variables in our models, FS/S and HLR, may or may not have changed over time. Investigation of this matter is left as a fruitful venue for future work. The fact that our DNN proxy remained accurate on both the 2004 BRFSS validation and partial validation from the 1995/2017 surveys suggests that the DNN proxy may have captured an intrinsic aspect of the relationship between the input and output variables at the state level, and may have some claim to generalizability. Further validation is warranted when more up-to-date BRFSS data on household gun ownership across all states is collected (as well as retraining and tuning of the DNN model). Third, our proxies were developed at the aggregate state level. They should not be used to infer gun ownership for other geographic areas or in subgroups of the population. Fourth, our proxies were developed using the average survey results from 2001 to 2002; as such, they are not suitable for use in longitudinal or time-series studies (see the discussions in the literature).[24,37]

Both of our new proxy measures provide meaningful accuracy improvements over the existing ones. They can be used for a variety of purposes in cross-sectional studies of firearm-related violence at the state level. Their two input variables are also easy to obtain, which makes the calculations simple and more convenient than accounting and controlling for a whole set of additional variables. To the best of our knowledge, this work provides the first use of DL analysis in this literature. We hope that the results will invite and encourage a broader adoption of these and other advanced ML methods in public health studies of firearm-related violence. We believe that significant opportunities reside in the adoption of ML in general, and DL in particular, for more accurate modeling of a host of public health problems, for teasing out novel insights from datasets, and ultimately for better-informed policy and decision making.

As a side note, a recent report by RAND, sponsored by Arnold Ventures, developed annual, state-level estimates of household gun ownership by combining data from surveys and proxy measures.[48] The report used multi-level regression with post-stratification to create state-level estimates of gun ownership from national surveys and then combined these estimates with proxy measures in a structural equation model to identify the latent gun ownership. The report suffers from some of the shortcomings identified in our work, for example, the failure to investigate statistical interaction and collinearity of the covariates or the failure to carefully examine non-linearities in the proxy-gun ownership relationships (e.g., the report notes, "the [hunting license] measure was heavily skewed to the right, so we performed a square root transformation"). The report also lacks model diagnostics and (performance) transparency. Although comparison of the RAND model with ours is not straightforward given the differences in approaches, we can compare the correlations of the estimated gun ownership rates in 2001, 2002, and 2004 with the BRFSS surveys. The RAND model, with four surveys and five proxy measures (male FS/S, female FS/S, hunting licenses, subscriptions to the *Guns & Ammo* magazine, and background checks—these are likely highly collinear), provided an average training R^2 of 97.5% over these years; our two-covariate DNN model provided a training R^2 of 98% in 2001–2002. Overfitting in the RAND model is therefore plausible, although it cannot be ascertained. No validation results were provided for further comparisons.

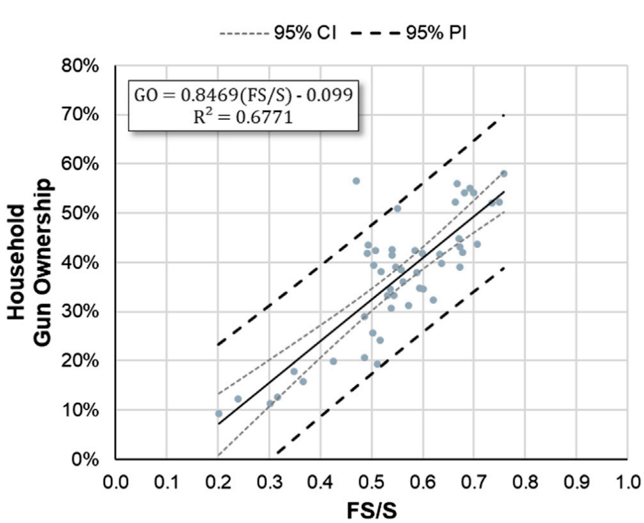

Figure 7. FS/S Proxy Measure (*Model 1*) of State-Level Household Gun Ownership with 95% Confidence and Prediction Intervals

EXPERIMENTAL PROCEDURES

Resource Availability

Lead Contact

David Gomez, M.S., dbgomez94@gmail.com.

Materials Availability

This study did not generate any materials.

Data and Code Availability

The DNN is freely available to be downloaded as a Python package from GitHub. Original data have been deposited to Mendeley Data: http://dx.doi.org/10.17632/bxsm39hsc9.1.

Data

State-Level Household Gun Ownership

We obtained state-level household gun ownership from the BRFSS,[22] which constitutes the gun ownership standard by which we trained and validated our new proxy measures. Although the BRFSS is the most extensive and direct measurement of state-level household gun ownership, the

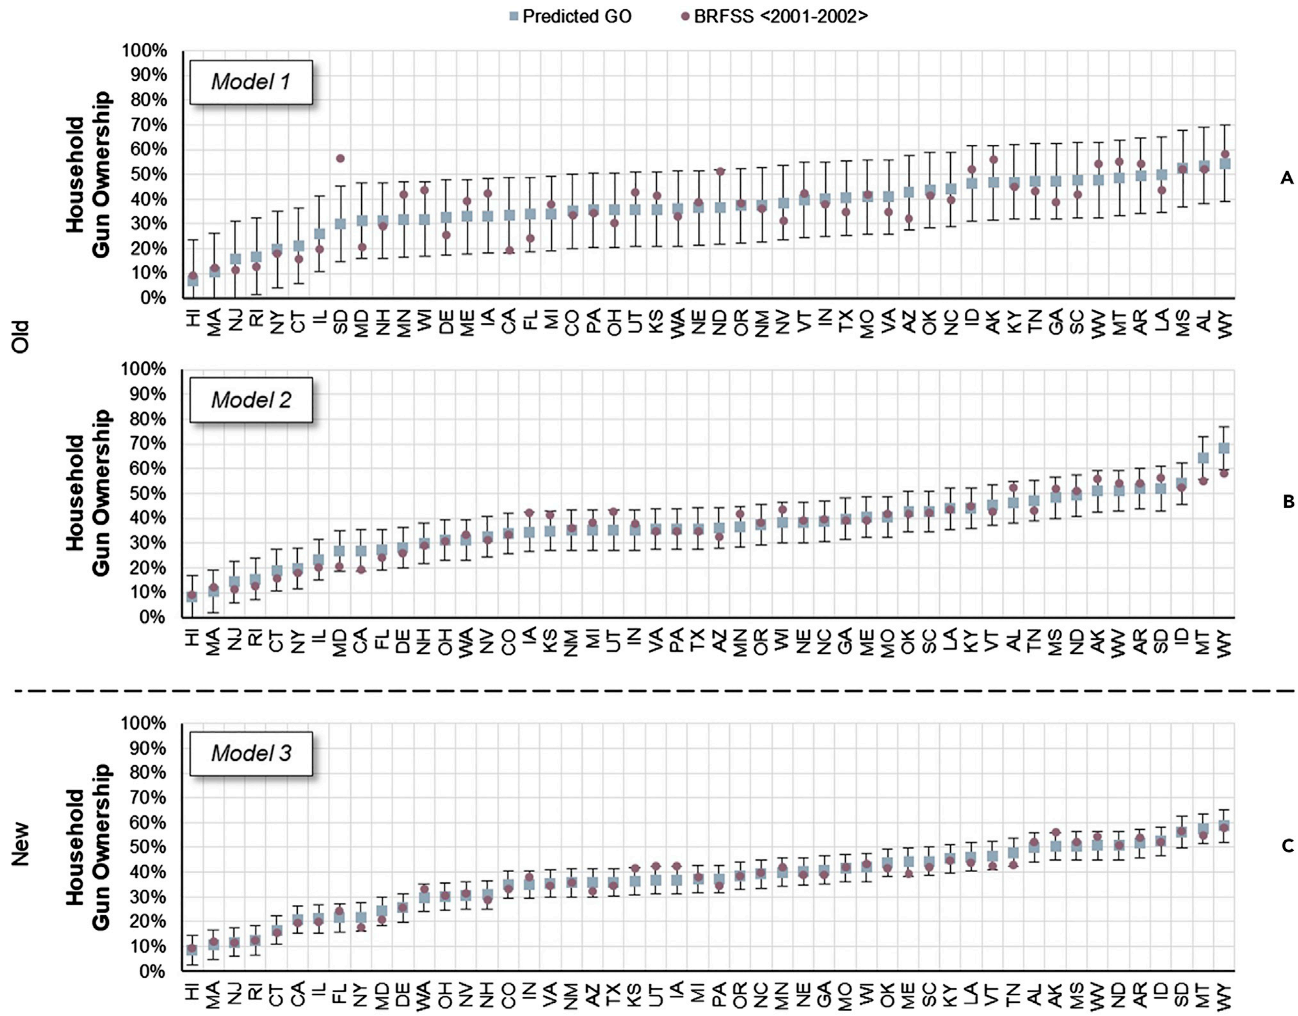

Figure 8. Comparison Between the Predicted State-Level Household Gun Ownership and that Estimated by the BRFSS (2001–2002 Average) for All Regression-Based Proxy Measures

Cook's Distances for (A) *Model 1*, (B) *Model 2*, and (C) *Model 3*. Error bars represent the 95% prediction intervals.

data have been limited over the years. The decision to include the gun ownership question, "Are any firearms kept in or around your home?" (or slight variation thereof) is made at the state level, and very often states do not include this question on their annual surveys. Shown in Table 3 are the available BRFSS survey data, along with the number of states which included the gun ownership question, as well as the number of observations.

To mitigate sampling variability in state-level household gun ownership without risking averaging out drift in the data, we chose to average the BRFSS figures over two consecutive years. The literature on gun ownership typically averages the BRFSS responses from 2001, 2002, and 2004, for which the gun ownership question was included.[37–39,49] From Table 3, we see that the BRFSS data from 2001 to 2002 represent the most comprehensive 2-year average, where 50 and 49 states included the gun ownership question on their surveys, respectively. We thus trained our models using the average 2001–2002 gun ownership data, and we use other survey years to validate our models. Note that California did not provide the gun ownership question in 2002; thus, the figure used was that of 2001. To adjust for the stratification, clustering, and sample weights of the BRFSS survey design, we used STATA's *svyset* function. The data quality and response rates are documented in the Summary Data Quality Reports.[50,51]

Portion of Suicides Committed with a Firearm

We obtained the portion of suicides committed with a firearm from the Centers for Disease Control and Prevention Web-based Injury Statistics Query and Reporting System (WISQAR).[52] In the training of our models, we also averaged the FS/S data over 2001–2002 to be consistent with our treatment of the BRFSS gun ownership data.

Hunting License Rates

We obtained state-level HLR from the US Fish and Wildlife Service.[53] As discussed in Siegel et al.,[37] this variable likely captures the gun ownership that FS/S is blind to, such as long rifles and shotguns, as handguns are predominantly used in suicides. We normalized the number of hunting licenses in a given state by the state's population aged 15+ and also averaged HLR over 2001–2002 in the training of our models.

Methods

Method 1: Regression Analysis

In developing our new regression-based proxy, *Model 3*, we examined different powers of FS/S and HLR and their statistical interaction to determine which combination provided the highest R^2 while maintaining statistical significance of the coefficients ($p < 0.05$). For thoroughness, we examined all combinations of powers from 0.1 to 5.0 in increments of 0.1. Note that the statistical interaction term was taken as the product of the transformed covariates. In this procedure, one combination of powers for each covariate and their interaction were trained on the dataset, and the R^2 and statistical significance of each coefficient were recorded. When all combinations were examined, the set of regressors that provided the highest R^2 while maintaining statistically significant coefficients was selected as our *Model 3* (see The New Regression-Based Proxy Measure).

Method 2: Deep Learning

In developing our new DL-based proxy, the DNN, we used the k-fold cross-validation technique to avoid overfitting and to assess how well the performance of the DNN will generalize to independent datasets (not used in the training of the DNN).[44,45] In k-fold cross-validation, the dataset is randomly split into k-subgroups or k-folds. Iterating through each fold, we set aside one fold as the "validation set" and the other $k - 1$ folds as the "training set." During each of the k-iterations, we optimized (trained) the DNN parameters on the training set then assessed its performance on the validation set. We then modified the DNN architecture, retrained the model, and repeated this assessment process on the other validation sets. In this work, we chose $k = 10$, such that for each step in this procedure, data for 45 states were used to train the DNN, and the remaining five states were used for validation. This value of k is typically used in practice because it provides a good balance in the bias-variance tradeoff.[54]

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers for their detailed comments and suggestions. We are indebted to the incredibly talented Ashling Devins for creating the graphical abstract. Support for the Open Access publication of this work was generously provided by the Injury Prevention Research Center at Emory University. This work is the result of an independent study and received no funding.

AUTHOR CONTRIBUTIONS

Conceptualization, D.B.G. and J.H.S.; Methodology, D.B.G., Z.X., and J.H.S.; Software, D.B.G. and Z.X.; Validation, D.B.G., Z.X., and J.H.S.; Formal Analysis, D.B.G., Z.X., and J.H.S.; Investigation, D.B.G., Z.X., and J.H.S.; Resources, J.H.S.; Data Curation, D.B.G.; Writing – Original Draft, D.B.G. and Z.X.; Writing – Review & Editing, D.B.G., Z.X., and J.H.S.; Visualization, D.B.G., Z.X., and J.H.S.; Supervision, J.H.S.; Project Administration, J.H.S.; Funding Acquisition, J.H.S.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: September 7, 2020
Revised: October 5, 2020
Accepted: November 3, 2020
Published: November 27, 2020

REFERENCES

1. National Research Council (2013). Priorities for Research to Reduce the Threat of Firearm-Related Violence (The National Academies Press). https://doi.org/10.17226/18319.
2. National Research Council (2005). Firearms and Violence: A Critical Review (The National Academies Press). https://doi.org/10.17226/10881.
3. 99th Congress (1985-1986). S.49—Firearms Owners' Protection Act. https://www.congress.gov/bill/99th-congress/senate-bill/49.
4. Lester, D. (1989). Gun ownership and suicide in the United States. Psychol. Med. *19*, 519–521.
5. Kellermann, A.L., Rivara, F.P., Somes, G., Reay, D.T., Francisco, J., Banton, J.G., Prodzinski, J., Fligner, C., and Hackman, B.B. (1992). Suicide in the home in relation to gun ownership. N. Engl. J. Med. *327*, 467–472.
6. Kellermann, A.L., Rivara, F.P., Rushforth, N.B., Banton, J.G., Reay, D.T., Francisco, J.T., Locci, A.B., Prodzinski, J., Hackman, B.B., and Somes, G. (1993). Gun ownership as a risk factor for homicide in the home. N. Engl. J. Med. *329*, 1084–1091.
7. Kaplan, M.S., and Geling, O. (1998). Firearm suicides and homicides in the United States: regional variations and patterns of gun ownership. Soc. Sci. Med. *46*, 1227–1233.
8. Kleck, G., and Hogan, M. (1999). National case-control study of homicide offending and gun ownership. Soc. Probl. *46*, 275–293.
9. Miller, M., and Hemenway, D. (1999). The relationship between firearms and suicide: a review of the literature. Aggression Violent Behav. *4*, 59–75.
10. Murnan, J., Dake, J.A., and Price, J.H. (2004). Association of selected risk factors with variation in child and adolescent firearm mortality by state. J. Sch. Health *74*, 335–340.
11. Shenassa, E.D. (2006). Utility of indices of gun availability in the community. J. Epidemiol. Commun. Health *60*, 44–49.
12. Miller, M., Hemenway, D., and Azrael, D. (2007). State-level homicide victimization rates in the US in relation to survey measures of household firearm ownership, 2001-2003. Soc. Sci. Med. *64*, 656–664.
13. Miller, M., Lippmann, S.J., Azrael, D., and Hemenway, D. (2007). Household firearm ownership and rates of suicide across the 50 United States. J. Trauma Inj. Infect. Crit. Care *62*, 1029–1035.
14. Miller, M., and Hemenway, D. (2008). Guns and suicide in the United States. N. Engl. J. Med. *359*, 989–991.
15. Wright, M.A., Wintemute, G.J., and Claire, B.E. (2008). Gun suicide by young people in California: descriptive epidemiology and gun ownership. J. Adolesc. Health *43*, 619–622.
16. Bangalore, S., and Messerli, F.H. (2013). Gun ownership and firearm-related deaths. Am. J. Med. *126*, 873–876.
17. Kposowa, A.J. (2013). Association of suicide rates, gun ownership, conservatism and individual suicide risk. Soc. Psychiatry Psychiatr. Epidemiol. *48*, 1467–1479.
18. van Kesteren, J.N. (2014). Revisiting the gun ownership and violence link: a multilevel analysis of victimization survey data. Br. J. Criminol. *54*, 53–72.
19. Anglemyer, A., Horvath, T., and Rutherford, G. (2014). The accessibility of firearms and risk for suicide and homicide victimization among household members. Ann. Intern. Med. *160*, 101–110.
20. Morrissey, T.W. (2017). Associations between active shooter incidents and gun ownership and storage among families with young children in the United States. Prev. Med. *100*, 50–55.
21. Anestis, M.D., and Houtsma, C. (2018). The association between gun ownership and statewide overall suicide rates. Suicide Life-Threatening Behav. *48*, 204–217.
22. Centers for Disease Control and Prevention (2001-2002). Behavioral Risk Factor Surveillance System Survey Data (US Department of Health and Human Services, Centers for Disease Control and Prevention).
23. Azrael, D., Cook, P.J., and Miller, M. (2004). State and local prevalence of firearms ownership measurement, structure, and trends. J. Quant. Criminol. *20*, 43–62.
24. Kleck, G. (2004). Measures of gun ownership levels for macro-level crime and violence research. J. Res. Crime Delinquency *41*, 3–36.
25. Miller, M., Azrael, D., and Barber, C. (2012). Suicide mortality in the United States: the importance of attending to method in understanding population-level disparities in the burden of suicide. Annu. Rev. Public Health *33*, 393–408.
26. Siegel, M., Ross, C.S., and King, C. (2013). The relationship between gun ownership and firearm homicide rates in the United States, 1981-2010. Am. J. Public Health *103*, 2098–2105.
27. Fleegler, E.W., Lee, L.K., Monuteaux, M.C., Hemenway, D., and Mannix, R. (2013). Firearm legislation and firearm-related fatalities in the United States. JAMA Intern. Med. *173*, 732–740.
28. Price, J.H., Thompson, A.J., and Dake, J.A. (2004). Factors associated with state variations in homicide, suicide, and unintentional firearm deaths. J. Community Health *29*, 271–283.

29. Rosenfeld, R., Baumer, E., and Messner, S.F. (2007). Social trust, firearm prevalence, and homicide. Ann. Epidemiol. *17*, 119–125.

30. Ruddell, R., and Mays, G.L. (2005). State background checks and firearms homicides. J. Crim. Just. *33*, 127–136.

31. Miller, M., Azrael, D., and Hemenway, D. (2002). Household firearm ownership and suicide rates in the United States. Epidemiology *13*, 517–524.

32. Miller, M., Azrael, D., and Hemenway, D. (2002). Rates of household firearm ownership and homicide across US regions and states, 1988-1997. Am. J. Public Health *92*, 1988–1993.

33. Miller, M. (2002). Firearm availability and suicide, homicide, and unintentional firearm deaths among women. J. Urban Health Bull. N.Y. Acad. Med. *79*, 26–38.

34. Miller, M., Azrael, D., and Hemenway, D. (2002). Firearm availability and unintentional firearm deaths, suicide, and homicide among 5-14 year olds. J. Trauma Acute Care Surg. *52*, 267–275.

35. Geier, D.A., Kern, J.K., and Geier, M.R. (2017). A longitudinal ecological study of household firearm ownership and firearm-related deaths in the United States from 1999 through 2014: a specific focus on gender, race, and geographic variables. Prev. Med. Rep. *6*, 329–335.

36. Reeping, P.M., Cerdá, M., Kalesan, B., Wiebe, D.J., Galea, S., and Branas, C.C. (2019). State gun laws, gun ownership, and mass shootings in the US: cross sectional time series. BMJ *364*, l542.

37. Siegel, M., Ross, C.S., and King, C. (2014). A new proxy measure for state-level gun ownership in studies of firearm injury prevention. Inj. Prev. *20*, 204–207.

38. Siegel, M., Negussie, Y., Vanture, S., Pleskunas, J., Ross, C.S., and King, C. (2014). The relationship between gun ownership and stranger and non-stranger firearm homicide rates in the United States, 1981–2010. Am. J. Public Health *104*, 1912–1919.

39. Siegel, M., Pahn, M., Xuan, Z., Fleegler, E., and Hemenway, D. (2019). The impact of state firearm laws on homicide and suicide deaths in the USA, 1991-2016: a panel study. J. Gen. Intern. Med. *34*, 2021–2028.

40. LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. Nature *521*, 436–444.

41. Maas, A.L., Hannun, A.Y., and Ng, A.Y. (2013). Rectifier nonlinearities improve neural network acoustic models. In Proceedings of the 30th International Conference on Machine Learning, *28*. http://robotics.stanford.edu/~amaas/papers/relu_hybrid_icml2013_final.pdf.

42. Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R.R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. ArXiv, 1207.0580 http://arxiv.org/abs/1207.0580.

43. Kingma, D.P., and Ba, J.A. (2019). ADAM: a method for stochastic optimization. ArXiv, 1412.6980 https://arxiv.org/abs/1412.6980v9.

44. Fushiki, T. (2011). Estimation of prediction error by using K-fold cross-validation. Stat. Comput. *21*, 137–146.

45. Kohavi, R. (1995). A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection. In International Joint Conference on Artificial Intelligence http://ai.stanford.edu/~ronnyk/accEst.pdf.

46. Weisberg, S. (2005). Applied Linear Regression, 3rd ed (Wiley-Interscience).

47. Hohman, F., Kahng, M., Pienta, R., and Chau, D.H. (2019). Visual analytics in deep learning: an interrogative survey for the next frontiers. IEEE Trans. Visual. Comput. Graph. *25*, 2674–2693.

48. Schell, T., Peterson, S., Vegetabile, B., Scherling, A., Smart, R., and Morral, A. (2020). State-Level Estimates of Household Firearm Ownership (RAND Corporation). https://www.rand.org/pubs/tools/TL354.html.

49. Briggs, J.T., and Tabarrok, A. (2014). Firearms and suicides in US states. Int. Rev. Law Econ. *37*, 180–188.

50. Centers for Disease Control and Prevention (2001). Summary Quality Data Report. https://www.cdc.gov/brfss/annual_data/annual_2001.htm.

51. Centers for Disease Control and Prevention (2002). Summary Quality Data Report. https://www.cdc.gov/brfss/annual_data/annual_2002.htm.

52. Centers for Disease Control and Prevention. WISQARS^{TM} (Web-Based Injury Statistics Query and Reporting System). https://www.cdc.gov/injury/wisqars/index.html.

53. US Fish & Wildlife Service. Historical Hunting License Data. https://www.fws.gov/wsfrprograms/Subpages/LicenseInfo/Hunting.htm.

54. James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). An Introduction to Statistical Learning (Springer).