

The social costs of gun ownership: a reply to Hayo, Neumeier, and Westphal

Philip J. Cook & Jens Ludwig

Empirical Economics
Journal of the Institute for Advanced
Studies, Vienna, Austria

ISSN 0377-7332
Volume 56
Number 1

Empir Econ (2019) 56:13-22
DOI 10.1007/s00181-018-1497-5



Your article is protected by copyright and all rights are held exclusively by Springer-Verlag GmbH Germany, part of Springer Nature. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".



The social costs of gun ownership: a reply to Hayo, Neumeier, and Westphal

Philip J. Cook^{1,2} · Jens Ludwig^{2,3}

Received: 14 November 2017 / Accepted: 23 November 2017 / Published online: 21 June 2018
© Springer-Verlag GmbH Germany, part of Springer Nature 2018

Abstract

We respond to the new article by Hayo, Neumeier, and Westphal (HNW), which is a critique of our 2006 article. The principal contribution of that article was to use a greatly improved proxy for gun prevalence to estimate the effect of gun prevalence on homicide rates. While the best available, our proxy, the ratio of firearms suicides to total suicides in a jurisdiction (FSS), is subject to measurement error which limits its use to larger jurisdictions that have enough suicides to stabilize the ratio. In this response, we report estimates for four different specifications and two data sets, the 200-county data and the data for the 50 states. We develop the claim that measurement error in FSS helps explain the observed pattern of results. Adopting the assumption that FSS follows a binomial process with a number of trials equal to the number of suicides, we characterize the relationship between measurement error and size of the jurisdiction, and thereby justify our conclusion that restricting the estimation to large jurisdictions reduces measurement error in FSS and hence the attenuation bias in the key coefficient estimate. We conclude that for the county-level data, the measurement error in FSS is of greater concern than using a specification that is flexible with respect to population. HNW focus on the latter but at the cost of increasing the effects of the former. We then demonstrate that the state-level data provide a robust case that more guns lead to more homicides.

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s00181-018-1497-5>) contains supplementary material, which is available to authorized users.

✉ Philip J. Cook
pcook@duke.edu
Jens Ludwig
jludwig@uchicago.edu

¹ Sanford School of Public Policy, Duke University, PO Box 90245, Durham, NC 27708, USA

² NBER, Cambridge, USA

³ School of Social Service Administration, University of Chicago, 969 East 60th Street, Chicago, IL 60637, USA

Keywords Measurement error · Crime · Guns · Homicide · Panel regression

JEL Classification C23 · C43 · K42

1 Introduction

In a new study, Hayo, Neumeier, and Westphal (hereafter HNW) revisit an article we published 12 years ago (Cook and Ludwig 2006a). Our analysis found evidence consistent with the hypothesis that more guns leads to more homicides—that is, both gun homicide rates and overall homicide rates tend to increase with rates of gun ownership at the county or state level. The goal of our 2006 paper, like that of all research, was to push forward the scientific frontier, not to be the last word on the topic. So we are grateful to HNW for working on the under-researched problem of gun violence in America, and for devoting so much thought to this problem.

HNW argue that our results depend on use of a problematic functional form in our regression analysis. They conclude that our paper provides little helpful information about the important public policy question of whether more guns lead to more, or less, homicide. We understand them to be making two points. First, we utilize homicide rates per capita as our dependent variable, rather than estimating a more flexible model that uses homicide counts as the dependent variable and population as an explanatory variable. They point out that our model, in log form, is equivalent to estimating their preferred model but with a restricted coefficient on log population (set to one in a model without additional covariates), and argue that it is preferable to use a more flexible form in which the data generate the coefficient estimate. Their estimates, based on their preferred approach, indicate a smaller and statistically insignificant effect of gun prevalence, although the estimated coefficient is still positive.

Second, in some of our regression specifications, we include covariates that, like our dependent variable, are normed to the population of the jurisdiction. HNW note that this specification is likely to produce biased coefficient estimates because of the so-called ratio fallacy. The use of a common denominator (population) in the dependent variable and several of the independent variables may indeed be problematic, especially if population is measured with substantial error. But as it turns out, the key differences between their results and ours emerge in the simple specification without covariates. For that reason, we do not pursue this aspect of their critique.

While technically correct, the first critique comes as something of a surprise since grouped-data regressions using this sort of specification are ubiquitous not just in the economics of crime but throughout all of applied economics.

In the next section, we briefly recap what we view as the primary contribution of our 2006 paper. The key challenges in the literature at the time our paper was written are related to the difficulties of measuring household gun prevalence, and to identifying the causal relationship between household gun prevalence and crime. The contribution of our paper relative to previous research was to use a greatly improved proxy for gun prevalence. (On the other hand, we made only modest progress improving on the identification strategies that have been used in past studies.) But while the best available, our proxy, the ratio of firearms suicides to total suicides in a jurisdiction

(FSS), is subject to measurement error which limits its use to larger jurisdictions that have enough suicides to stabilize the ratio.

In Sect. 3, we characterize HNW's critique in more detail. In Sect. 4, we report estimates for four different specifications and two data sets, the 200-county data and the data for the 50 states. We develop the notion that measurement error in FSS helps explain the observed pattern of results. Adopting the assumption that FSS follows a binomial process with a number of trials equal to the number of suicides, it is possible to characterize the relationship between measurement error and size of the jurisdiction, and thereby justify our conclusion that restricting the estimation to large jurisdictions reduces measurement error in FSS and hence bias in the key coefficient estimate.

We conclude that for the county-level data, the measurement error in FSS is of greater concern than using a specification that is flexible with respect to population. HNW focus on the latter but at the cost of increasing the effects of the former. We go on to conclude that the state-level data provide a robust case that more guns lead to more homicides. We recognize that there is room for future studies to improve on ours with respect to identifying the causal effect of gun prevalence.

2 The contribution of Cook and Ludwig (2006a)

Prior to the publication of our 2006 paper, there were two prominent studies that came to qualitatively different conclusions about the relationship between prevalence of gun ownership and crime (Lott 2000; Duggan 2001).¹ We hypothesized that one important difference between these studies was in how they measured gun prevalence. The challenge arises because the USA does not require gun owners to register their firearms with the government, so there are no administrative records that document the number of guns in private hands within an area or the share of households that keep guns.

Lott (2000) carried out what is essentially a cross-section analysis using state-level data, measuring gun prevalence from voter exit polls in 1988 and 1996 that happened to include an item on gun ownership.² He estimated the elasticity of homicide with respect to gun ownership rates to be equal to -3.3 . But of course those who choose to vote in any given election are a highly unrepresentative sample of a jurisdiction's population, and the characteristics of voters tend to change from election to election depending on turnout and other factors. Furthermore, responses to survey questions about gun ownership distort the truth in somewhat predictable ways (see, for example, Ludwig et al. 1998). As we noted in our 2006 paper, some indication of the problem with Lott's proxy measure comes from the fact that it suggests that gun ownership rates increased substantially across the USA over the period from 1988 to 1996, from 27.4 to 37.0%. Yet the best available source of national data on how gun ownership

¹ For a broader discussion of the scientific evidence about guns and crime, such as what is known about whether guns increase the lethality of violence and related issues, see Cook and Ludwig (2006b) and Cook and Goss (2014).

² As we noted in our 2006 paper, Lott pools his data from two waves of voter exit polls, in 1988 and 1996, but in his regressions controls for region but not state fixed effects. So most of the variation in gun ownership rates in his analysis will be from across-state variation, rather than within-state over-time variation. .

rates have changed over time in the USA, from the General Social Survey, indicates that individual gun ownership rates changed very little over this period (see Kleck 1997, pp. 98–99; Smith and Son 2015).

Duggan (2001) assembled panel datasets of both state- and county-level information, using as his proxy for local gun ownership prevalence the subscription rate to *Guns and Ammo* magazine. In contrast to Lott's estimates, which indicate that more guns in an area generate a large reduction in homicides, Duggan's findings indicate that more guns increase homicides, with an elasticity estimate of +.2.

Perhaps the main contribution of our own study was to use what we believe is the most reliable proxy currently available for gun ownership at the local and state level in the USA: the fraction of suicides that are committed with a firearm, FSS. Two independent studies validated FSS as a remarkably accurate proxy for gun prevalence across jurisdictions (Azrael et al. 2004; Kleck 2004), and concluded that it was the best of all available proxies that had been tried up until that point. In our 2006 paper, we reported new evidence in support of the validity of FSS in measuring changes in gun prevalence over time.

One limitation of the FSS proxy is that suicide is a rare event (with an annual rate on the order of 1 in 10,000), so that it is unstable for smaller jurisdictions. That fact creates a tradeoff for the analyst; for example, we report results for the 200 largest counties in the USA, but also the 100 largest and 50 largest—and also at the state level. The smaller data sets have less statistical power, but eliminating smaller jurisdictions means that FSS has less random variability than in the 200-county data set. We return to this issue below.

Using a standard panel-data setup with county and year fixed effects, we estimate an elasticity of homicide rates to FSS (lagged 1 year) of .100, with standard error of .044 without covariates, and .086 (.038) with a set of covariates that includes the log of burglary and robbery rates and socio-demographic characteristics of the county. Also, as a point of clarification: contrary to what HNW report, the standard errors included in our 2006 paper were indeed clustered at the county level in this analysis.

To reduce measurement error in FSS, we recalculated our estimates using a 2-year average of FSS within a county; as noted above we also tried restricting our sample to the largest 100 and then the largest 50 counties, and finally using state-level instead of county-level data. Consistent with the idea that increasing the number of suicides used to calculate FSS reduces measurement error, and hence attenuation bias, our estimated elasticities are larger for the samples that are restricted to larger jurisdictions, especially for the state-level analysis. Incidentally, HNW do not attempt to replicate these models or explore their sensitivity to changing the functional form assumption they are concerned about.

In our view, the main limitation of our analysis was a lack of a clearly exogenous source of variation in gun prevalence. Duggan (2001) relied on standard panel-data approaches to deal with this problem. Our study made only modest progress on that challenge. We used the 1-year lag of FSS to establish that changes in gun prevalence preceded changes in homicide rate, and we demonstrated that the results were robust to including census division-by-year fixed effects.

3 HNW's critique

The new paper by HNW raises a very different concern about this literature that, as far as we know, has not come up before. The concern pertains to a fairly subtle functional form assumption implicit in the use of homicide and other variables in per capita form. They note that our specification restricts the way that the population of an area is allowed to affect the number of homicides that occur, and report results suggesting that estimates based on a less restrictive specification are much weaker. Before engaging this point directly, we first discuss the disagreement between their results and ours for our original specification.

When they attempt to estimate our primary specification using 200 counties, they find smaller (though still significant) positive coefficients for the relationship between log-lagged FSS and log homicide rate. In preparing this response, we re-ran these regressions and reproduced our original results exactly. Apparently HNW are using a somewhat different data set than are we.

We utilized the homicide and suicide counts from the National Vital Statistics program. These data tend to be of high quality—and the homicide data, in particular, are more reliable than those from the FBI's Uniform Crime Reports—but still have a few large errors (Loftin et al. 2008). Our original paper included an appendix of Electronic Supplementary Material explaining the choices we made in cleaning these data. We were unable to determine what source of homicide data HNW used, or how they cleaned it.

What we take to be their main critique can be quickly summarized. Our principal specification in simplest form (without covariates) was as follows:

$$\text{Log}(H_{it}/N_{it}) = \beta_0 + \beta_1 \text{LogFSS}_{it-1} + d_i + d_t + \varepsilon_{it} \quad (1)$$

In this equation, H_{it} is the number of homicides that occur in jurisdiction (i) in year (t), N_{it} is the population in that jurisdiction, FSS_{it-1} is the share of suicides committed with a firearm the previous period (a proxy for the prevalence of gun ownership), and d_i and d_t are jurisdiction and year fixed effects, respectively.

A variation on Eq. (1) is Eq. (2), where the denominator of the dependent variable is now on the right-hand side as a covariate.

$$\text{Log}(H_{it}) = \beta_0 + \beta_1 \text{LogFSS}_{it-1} + \beta_2 \log(N_{it}) + d_i + d_t + v_{it} \quad (2)$$

HNW point out that our model is equivalent to estimating Eq. (2) with a restriction that $\beta_2 = 1$ (although they fail to point out that the structure of the error term is entirely different in these two versions). Put differently, our specification implicitly imposes the restriction that population change is by itself neutral with respect to the homicide rate. A (say) .1 log point change in population within an area translates into a .1 log point change in homicides. As HNW note, in principle one could estimate Eq. (2) instead and treat β_2 as a parameter to be estimated from the available data.

That the implicit restriction we imposed on β_2 could have important implications for our estimates is indeed not something we had considered previously, partly because

this restriction is standard practice within the economics of crime literature,³ and not just for that literature. Indeed, this approach is the norm for empirical papers within applied economics, including numerous papers that have become citation classics. To explore this point, we scanned the 100 most-cited economics papers published between 1970 and 2006, finding 14 that included regression results with grouped data (such as jurisdiction-level rates) as the unit of observation.⁴ All 14 of these specified the regressions with a rate per capita as dependent variable, while not including population as a covariate. Could it really be that there is a fundamental problem working in rates per capita that afflicts not just our study but much of the empirical literature in economics?

Finally, we note that they also challenge our proxy FSS on the grounds that the denominator of FSS (the number of suicides) is correlated with population. In their preferred specification, they estimate the effects of gun suicides and suicides separately. But that approach fails to recognize that it is the ratio, FSS, that has been validated as a proxy for gun prevalence, not the numerator and denominator separately. In any case, HNW themselves report that the results are not sensitive to the choice of whether to use FSS as the key explanatory variable of interest versus including the number of gun suicides and total suicides as separate explanatory variables.

4 New estimates

In Table 1, we report a new set of findings that is responsive to HNW's critique. Each cell is the FSS coefficient estimate from a different regression. The first row utilizes the 200-county data set. In that row, the first cell reproduces a result from our original paper, based on Eq. 1. The other three cells are based on Eq. 2, with the homicide count (rather than rate) as the dependent variable, either logged (in the OLS form) or the equivalent in a negative binomial specification. In accord with HNW, we find that the key coefficients for Eq. 2 are smaller than for Eq. 1 and statistically insignificant.

The second row reports the same estimates but now for our state-level data set. Again, the first cell reproduces our original result (one that is not considered by HNW). The other three estimates based on state-level data are very similar to the original in this case, a robust set of results that indicate a far larger influence of gun prevalence on homicide than suggested by the county data. For the state-level data, it really does not matter whether population is included as a covariate.

One explanation for the pattern of results in Table 1 is in terms of measurement error of FSS. In a simple regression, measurement error in the explanatory variable attenuates the coefficient estimate toward zero. We would expect that FSS will be subject to more random variation in smaller jurisdictions, so that county-level data

³ Crime rates per capita is the dependent variable of interest not just for Lott (2000) and Duggan (2001), but also Levitt's (1996) study of prisons and crime, Levitt's (1997) study of the effects of additional police on crime, which was subject to a careful re-analysis by McCrary (2002) with no mention of any concerns about working in crime rates, Raphael and Winter-Ebmer's (2001) study of the link between unemployment rates and crime, and Donohue and Levitt's (2001) study of the effects of legalized abortion on crime, which has been the subject of intensive public scrutiny almost unheard of in academic research.

⁴ We utilized the list of most cited articles in economics between 1970 and 2005, taken from Kim et al. (2006).

Table 1 The effect of gun prevalence on homicide rates: CL versus HNW

Model specification/data	OLS using CL specification ^a	Negative binomial regression with CL constraint on population variable ^b	OLS using HNW specification ^c	Negative binomial regression version of HNW specification ^d
Panel regression results for annual data, 1980–1999 ^{a,e}				
200 largest counties	.100 ** (.044)	.076 (.069)	.037 (.049)	.014 (.060)
50 states ^e	.588** (.237)	.677** (.265)	.546** (.197)	.614*** (.213)

*, **, and *** Significance at the .10, .05, and .01 levels, respectively

^aOLS regression: Log homicide rate regressed on log lag FSS and county (or state) fixed effects and year fixed effects, but no other covariates, weighted by population. SEs are cluster corrected

^bNegative binomial regression: log homicide count regressed on log lag FSS, log population, and county (or state) fixed effects and year fixed effects, weighted by population. The coefficient on log population is constrained to be 1

^cOLS regression: same as b, but with coefficient on log population unconstrained

^dNegative binomial regression: same as c. Instead of using population as an exposure variable, we include log population as a right-hand side variable

^eEstimates based on data for 1982–1999

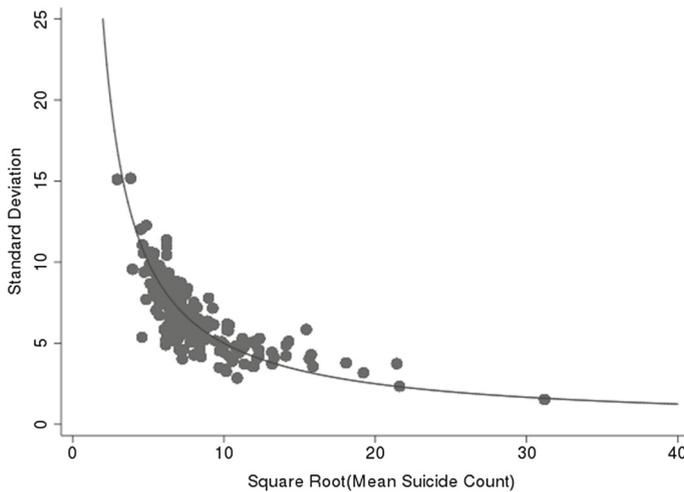


Fig. 1 Observed and theoretical measurement error in FSS, 200 counties, 1980–2000

would have greater measurement error than state-level data. This intuition can be made more precise. To model the error, we assume that the choice of weapon for each suicide in jurisdiction i , year t , is an independent trial with probability p_{it} of being a gun. The standard deviation for any one year/jurisdiction is the square root of $p(1-p)/S$, which is well approximated by $.5/\text{Sq Rt}(S)$ for values of p in the range $.4-.6$. This model fits quite well, as suggested by Fig. 1, which plots the observed standard deviations of FSS for each jurisdiction over the 20-year span against the square root of the mean number of suicides for that jurisdiction (after detrending FSS). The superimposed line is the theoretical relationship based on the binomial model.

Note that the population variable itself may also be measured with error. Over the time period of our 2006 study the primary source of information about county or state populations came from the decennial census carried out by the US Census Bureau. The challenge becomes measuring population during the inter-censal years; while vital statistics information can capture year-to-year population changes within an area due to births and deaths, there is no ideal way to capture in- or out-migration.

The effect of the HNW specification, then, is to run a regression with two explanatory variables that are both measured with error: FSS and population. The exact relationship between the errors of measurement in these two variables is not obvious, but there is no reason to believe they are independent. That may further increase the attenuation bias in the coefficient on FSS, although that is not necessarily the case. In fact, we cannot predict the sign, much less the magnitude, of the biases that result for the coefficients on either FSS or population (see for example Cameron and Trivedi 2005).

This explanation for the problems introduced by HNW's specification also helps explain why the state-level regression results do not change nearly as much as the county-level analysis when run using the HNW specification rather than our original specification as in our 2006 paper. Figure 2 suggests that the measurement error in

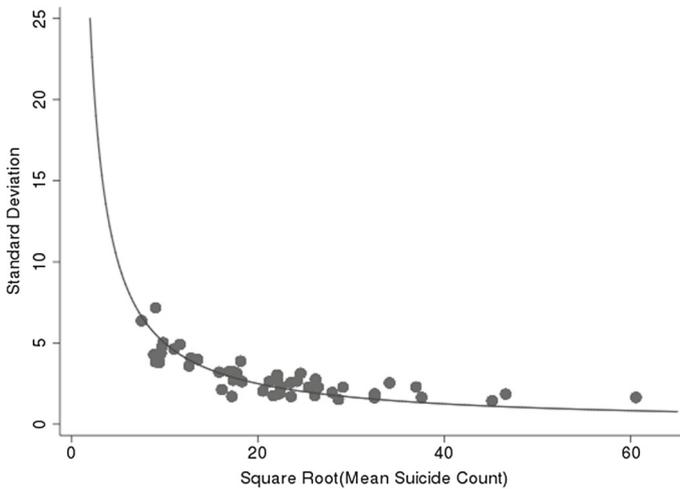


Fig. 2 Observed and theoretical measurement error in FSS, 50 states

FSS is much less pronounced compared to what we see for the county-level data shown in Fig. 1. In addition, we can assume measurement error is less pronounced for population at the state than county level, partly because migration across county lines will be more common than will be moves across state lines.⁵

5 Concluding thoughts

HNW have posed two unusual challenges, not only to our 2006 paper, but also to much of the literature in empirical economics. Their primary challenge has motivated our close inquiry into measurement error in FSS. The result is to increase our confidence in the very positive results we found using these data in our 2006 paper, suggesting that the negative externality of gun ownership is much higher than estimated from the 200-county data.

Acknowledgments Thanks to Lauren Spiegel for outstanding assistance with the data analysis, to Dan Black and Seth Sanders for helpful conversations. Any errors and all opinions are of course our own.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval This article does not contain any studies with human participants or animals performed by any of the authors.

⁵ See for example: <http://www.pewsocialtrends.org/files/2010/10/Movers-and-Stayers.pdf>.

References

- Azrael D, Cook PJ, Miller M (2004) State and local prevalence of firearms ownership measurement, structure, and trends. *J Quant Criminol* 20(1):43–62
- Cameron AC, Trivedi PK (2005) *Microeconometrics: methods and applications*. Cambridge University Press, New York
- Cook PJ, Goss KA (2014) *The gun debate: what everyone needs to know*. Oxford University Press, New York
- Cook PJ, Ludwig J (2006a) The social costs of gun ownership. *J Public Econ* 90:379–391
- Cook PJ, Ludwig J (2006b) Aiming for evidence-based gun policy. *J Policy Anal Manag* 25(3):691–735
- Donohue JJ, Levitt SD (2001) The impact of legalized abortion on crime. *Q J Econ* 116(2):379–420
- Duggan Mark (2001) More guns, more crime. *J Polit Econ* 109(5):1086–1114
- Hayo B, Neumeier F, Westphal C (forthcoming) The social costs of gun ownership revisited. *Empir Econ*
- Kim EH, Morse A, Zingales L (2006) What has mattered to economics since 1970. *J Econ Perspect* 20(4):189–202
- Kleck G (1997) *Targeting guns: firearms and their control*. Aldine de Gruyter, New York
- Kleck G (2004) Measures of gun ownership levels for macro-level crime and violence research. *J Res Crime Delinq* 41(1):3–36
- Levitt SD (1996) The effect of prison population size on crime rates: evidence from prison overcrowding litigation. *Q J Econ* 111(2):319–351
- Levitt SD (1997) Using electoral cycles in police hiring to estimate the effect of police on crime. *Am Econ Rev* 87(3):270–290
- Loftin C, McDowall D, Fetzer MD (2008) A comparison of SHR and Vital Statistics Homicide estimates for US cities. *J Contemp Crim Justice* 24(1):4–17
- Lott JR (2000) *More guns, less crime*, 2nd edn. University of Chicago Press, Chicago
- Ludwig J, Cook PJ, Smith TW (1998) The gender gap in reporting household gun ownership. *Am J Public Health* 88(11):1715–1718
- McCrary J (2002) Using electoral cycles in police hiring to estimate the effect of police on crime: comment. *Am Econ Rev* 92(4):1236–1243
- McDowall D, Loftin C, Fetzer MD (2008) A comparison of SHR and vital statistics homicide estimates for U.S. cities. *J Contemp Crim Justice* 24(1):4–17
- Raphael S, Winter-Ebmer R (2001) Identifying the effect of unemployment on crime. *J Law Econ* 44(1):259–283
- Smith TW, Son J (2015) Trends in gun ownership in the United States, 1972–2014 (General Social Survey Final Report). NORC at the University of Chicago, Chicago. http://www.norc.org/PDFs/GSS%20Reports/GSS_Trends%20in%20Gun%20Ownership_US_1972-2014.pdf. Downloaded 28 Oct 2017