Journal of Public Economics 90 (2006) 379–391
JOURNAL OF PUBLIC ECONOMICS
www.elsevier.com/locate/econbase

The social costs of gun ownership

Philip J. Cook[a,*], Jens Ludwig[b]

[a] *Duke University and NBER, United States*
[b] *Georgetown University and NBER, United States*

Received 17 August 2004; received in revised form 18 February 2005; accepted 18 February 2005
Available online 31 May 2005

Abstract

This paper provides new estimates of the effect of household gun prevalence on homicide rates, and infers the marginal external cost of handgun ownership. The estimates utilize a superior proxy for gun prevalence, the percentage of suicides committed with a gun, which we validate. Using county- and state-level panels for 20 years, we estimate the elasticity of homicide with respect to gun prevalence as between +0.1 and +0.3. All of the effect of gun prevalence is on gun homicide rates. Under certain reasonable assumptions, the average annual marginal social cost of household gun ownership is in the range \$100 to \$1800.

Jel classifications: H21; I18; K42
Keywords: Regulation; Social costs; Evaluation methods; Externalities; Gun violence

1. Introduction

Like many other private decisions about health and safety, such as getting vaccinated, purchasing LoJack (Ayres and Levitt, 1998), or driving a sport utility vehicle (Gayer, 2004), private gun ownership may impose externalities. Widespread gun ownership in a community could provide a general deterrent to criminal predation, lowering the risk to owners and non-owners alike. But widespread gun ownership could also lead to increased

* Corresponding author.
E-mail address: pcook@duke.edu (P.J. Cook).

0047-2727/\$ - see front matter
doi:10.1016/j.jpubeco.2005.02.003

risks of various sorts, including the possibility that guns will be misused by the owners or transferred to dangerous people through theft or unregulated sale. Whether the social costs of gun ownership are positive or negative is arguably the most fundamental question for the regulation of firearms in the United States.

Previous research has produced conflicting conclusions. One prominent estimate for the effects of gun prevalence on homicide is by John Lott (2000), who relates state-level estimates of gun ownership rates from voter exit polls in 1988 and 1996 to state crime rates, conditioning on several socioeconomic variables in a cross-section analysis. His estimate of the elasticity of homicide with respect to state gun ownership rates is extraordinarily large, equal to -3.3.[1]

Mark Duggan (2001) identifies the relationship between guns and crime using over-time variation in panels of states and also counties. Duggan's elasticity estimate, +0.2, is of the opposite sign from Lott's and an order of magnitude smaller. There is some question about the validity of his proxy for gun prevalence, the subscription rate to *Guns and Ammo* magazine.

In this paper we follow Duggan's lead in using panel regression methods to estimate the effect of gun prevalence on homicide rates, but with a different and well-validated proxy variable. Our results suggest that the social cost of an additional household acquiring a handgun depends on the rate of violence and the existing prevalence of guns, but under a wide range of assumptions is greater than $100 per year.

2. FSS as a proxy for gun prevalence

Since most states lack any sort of registration or licensing system that would generate administrative data on firearms ownership, household surveys provide the only direct source of information on this matter. But survey data are not always available or reliable for sub-national units, so analysts have employed a variety of proxy variables. Two independent inquiries have recently identified one such proxy as superior to all others for the purpose of estimating the cross-section structure of gun prevalence across large geographic entities (Azrael et al., 2004; Kleck, 2004). That proxy is the fraction of suicides committed with a firearm (FSS).

Our use of FSS is primarily to estimate variation over time rather than in the cross-section. To validate this use requires consistent estimates of gun prevalence over time, preferably at a sub-national level. The "gold standard" for national surveys of gun ownership is the General Social Survey (GSS). We ran panel regressions of GSS-based estimates of gun prevalence against two proxies, FSS and the subscription rate to *Guns and Ammo*, the proxy used by Duggan (2001). The estimated coefficients of our GSS

[1] Lott conditions on region but not state dummies in his regressions, so his estimates will be identified primarily by cross-sectional variation in gun ownership rates (Azrael et al., 2004). A more fundamental problem is that there are serious problems with his voter exit poll data, which suggest that from 1988 to 1996 gun ownership rates increased for the U.S. as a whole from 27.4 to 37.0% (p. 36). Yet the best source of national data on gun ownership trends – the General Social Survey – indicates that individual gun ownership trends were essentially flat during this period (Kleck, 1997, pp. 98–99).

measures on FSS are in every case significantly positive, and are especially strong when year fixed effects are omitted, while the subscription rate to *Guns and Ammo* performs less well and in some cases yields a negative coefficient estimate.[2]

3. Data

The estimates presented below are based on panel data for 200 counties that had the largest population in 1990,[3] or a subset of those counties, for the period 1980 to 1999. We also present estimates based on state-level panel data. The 200 largest counties accounted for 74% of all homicides in the United States in 1990.[4]

Suicide and homicide counts are taken from Vital Statistics Program mortality data, based on reports of coroners and medical examiners and compiled by the National Center for Health Statistics. Data on robbery, burglary, and other types of crime besides homicide are from the FBI's Uniform Crime Reports.

Finally, we control for other changes over time in county socio-demographic characteristics that could affect both crime and gun prevalence. Such data are available at the county level only from the decennial Census, from which we interpolate data for the inter-Censal years. Covariates include the prevalence of blacks, households headed by a female, urban residents, and residents living in the same house 5 years ago.

4. Empirical strategy

The basic empirical approach here is to estimate the relationship between gun prevalence and homicide by exploiting the substantial across-area differences in trends in gun ownership over a 20-year period. Our baseline estimates are generated from model (1), which relates the natural log of jurisdiction (i)'s homicide rate (or, alternatively, the gun- or non-gun homicide rate) in year t against FSS, the proxy for the jurisdiction's gun ownership rate, in year $(t-1)$. FSS is lagged by one period out of concern for reverse causation – gun ownership may be consequence as well as cause of a county's crime rate – although the lag can also be justified for substantive reasons: the thefts and secondary-

[2] The GSS is conducted by the National Opinion Research Center most years from 1972 to 1993 and biennially since 1994 (Davis and Smith, 1998), and is capable of providing representative samples at the national or census region or even division level. Our panel dataset for this validation exercise is defined over the nine Census divisions and the 14 years in which GSS fielded gun questions between 1980 and 1998. In these regressions, we condition on fixed effects for Census division in all model specifications. We define "prevalence" in the GSS data for either handguns or all guns, and for either households or individuals. For additional details, see Cook and Ludwig (2004b).

[3] Kelly (2000) used this sample of counties in studying the determinants of crime rates. The 5 counties of New York City are combined in our analysis due to data limitations. Oklahoma City was dropped in 1995 due to the large homicide count associated with the bombing of the federal building there.

[4] Also of some interest is what fraction of all guns in the U.S. is found in the top 200 counties. While we cannot perform this calculation with our FSS proxy, which is not available for all counties, we find that 43% of all *Guns and Ammo* subscriptions in the U.S. are in the 200 largest counties.

market transfers that move guns from households to use by criminals will ordinarily take some time. To further control for the possibility of reverse causation, we condition on the natural log of the area's burglary and robbery rates, which are the kinds of crimes that seem likely to motivate the acquisition of a firearm for self-defense. These crime variables also are a good reflection of criminogenic factors in the community that influence homicide rates (Blumstein, 2000). To account for other county or state characteristics that affect homicide, the regression model includes year and county/state fixed effects, as well as the logs of the socio-demographic variables. The regression estimates are weighted by each county or state's population to account for heteroskedasticity in the error term.

$$\log\, Y_{it} = \beta_0 + \beta_1 \log \mathrm{FSS}_{it-1} + \beta_2 X_{it} + d_i + d_t + \varepsilon_{it}. \tag{1}$$

Another concern is serial correlation in the error structure, given that FSS changes only slowly over time within counties and that other unmeasured determinants of county crime rates may also have jurisdiction-specific trends.[5] We address this problem by calculating Huber–White standard errors that are robust to an arbitrary autocorrelation pattern in the errors over time within counties. Bertrand et al. (2004) show that this approach works better than more parametric strategies in panels with a short time dimension.

A final concern in estimating Eq. (1) is that the proxy for gun prevalence, FSS, is subject to measurement error of two types. First, because it is only a proxy, the correlation between FSS and the "true" prevalence is presumably less than one. Judging the quality of the proxy in that sense is difficult, given that there are no error-free measures of the criterion variable. In particular, survey-based estimates are subject to sampling error and reporting error. Based on an analysis of national GSS estimates over time, the hypothesis that FSS *is* a "perfect" proxy cannot be rejected, but that is not the same thing as demonstrating that it is perfect in fact.[6]

Second, and probably more important, is that the reliability of FSS will depend on the number of suicides used to compute it. For the 21 years of data on 200 large counties, the 10th and 90th percentiles have 27 and 142 suicides respectively, with a median of 52 and a mean of 196. If the choice of weapon in suicide follows a binomial process, then a jurisdiction with 50 suicides a year would generate an observed FSS that is subject to a standard error of 7 percentage points. The effect of this measurement error will be to bias the coefficient estimate of FSS toward zero. We address this problem in a variety of ways below, including re-calculating our estimates with state-level data. While the state data have the advantage of reducing measurement error in FSS, one drawback is that county-

[5] Testing for the presence of serial correlation in fixed-effects models is complicated in applications where the time dimension is fairly short compared to the number of observational units. Following Solon (1984), we test for serial correlation by first-differencing the data, and then keep the residuals from a regression of the log change in homicides against the log change in FSS and year effects. A regression of these residuals against their 1-year lag yields a coefficient of -0.4, close to the value of -0.5 characteristic of an error structure that is serially uncorrelated. Additional tests indicate that serial correlation is a somewhat greater problem with the state-level data.

[6] The correlation between national household handgun prevalence and FSS over 18 waves of the GSS is 0.635, very close to the mean of a large number of correlations generated from a simulation based on the assumption that FSS is exact and the GSS estimates are unbiased but subject to normal sampling error. That mean is 0.664.

Table 1
Descriptive statistics for county data

	Full sample (largest 200)	Bottom quartile 1980 FSS	Top quartile 1980 FSS
Full period (1980–1999)			
FSS	49.9	34.6	66.9
Homicide rate	11.0	10.9	14.4
Gun homicide rate	7.3	6.9	10.1
%Urban	92.6	94.7	91.8
%Percent black	14.0	13.5	19.5
%Female household head	18.0	20.1	18.5
# Suicides	195.8	192.5	120.0
FSS in selected years			
1980	48.0	29.2	73.3
1990	52.8	37.2	69.1
1999	48.0	34.9	59.8

Source: Mortality—National Center for Health Statistics, Vital Statistics, Mortality; Crime—Federal Bureau of Investigation, Uniform Crime Reports; Demographics—US Census Bureau, 1980, 1990 and 2000 Censuses.

level gun prevalence may be more relevant for local gun availability in the used or "secondary" gun market (Cook et al., 1995).[7]

5. Results

Table 1 presents descriptive statistics for the full panel assembled from annual data for the 200 largest counties for the years 1980–1999 (all calculations are weighted by county population). Over the entire sample period, the average homicide rate is 11 per 100,000 residents, with half of all suicides having been committed with a firearm.

Table 1 also provides some sense for the variation in gun ownership that identifies the panel data estimates shown below. The second and third columns of Table 1 present data for the top and bottom quartiles for our 200 counties ranked according to their gun ownership rates at the start of our panel, in 1980. The (disproportionately Southern) counties where guns are most common in 1980 experience a persistent and pronounced reduction in household gun ownership rates during the 20 years of our panel, as reflected by the nearly 20% decline in FSS over this period. At the same time, counties where guns were least common in 1980 (disproportionately in the Northeast and Midwest regions) experienced an increase in FSS of 20% from 1980 to 1999.

The source of this convergence remains something of a mystery (Azrael et al., 2004). If whatever drove this convergence between high- and low-gun ownership areas was orthogonal to the determinants of homicide trends, then a difference-in-differences

[7] On the other hand, a potential advantage of the state-level data comes from the possibility that people cross county lines to obtain firearms. This may not be a very severe problem, at least for youth, who account for a disproportionate share of all gun crime. When Cook and Ludwig (2004a) regress an indicator for youth gun carrying against FSS the relationship is much stronger when FSS is measured at the county than at the state level.

estimate of the effect of FSS on homicide (Y) would be unbiased. In particular, expression (2) is an estimate of the elasticity of Y with respect to FSS, where Δ indicates the difference between 1999 and 1980, and the subscripts $Q1$ and $Q4$ refer to "top quartile" and "bottom quartile" respectively.

$$\left(\Delta\ln Y_{Q1} - \Delta\ln Y_{Q4}\right) / \left(\Delta\ln\mathrm{FSS}_{Q1} - \Delta\ln\mathrm{FSS}_{Q4}\right). \tag{2}$$

The elasticity of homicide with respect to FSS estimated in this fashion is +0.18. A similar calculation for gun homicides yields an elasticity with respect to gun ownership rates of +0.35. These simple estimates turn out to be quite compatible with those derived from the panel regression analysis that uses all of the variation across counties over time.

5.1. Panel regression findings

The first column of Table 2 presents the results for our most parsimonious model, which includes county and year fixed effects but no other covariates. The estimated elasticity of homicide with respect to the lagged value of log FSS equals +0.100 ($p<0.05$). The final three columns of Table 2 show that this point estimate is not sensitive to controlling for several sets of influential covariates.

Table 3 reports results for a number of alternative specifications, in each case for three dependent variables: the logs of the homicide rate, the gun homicide rate, and the non-gun homicide rate. If the predominant causal mechanism linking gun prevalence to homicide is that increased prevalence induces substitution of guns for other weapons in assaults, with a consequent increase in lethality, then only the gun homicide rate will increase in response to an increase in FSS. Table 3 generally supports this prediction.

The results are robust to a variety of modifications to our basic estimation approach. In the second row, additional county-level characteristics are added –percentages of resident

Table 2
Baseline results, county-level data, 1980–1999

	Ln(Hom)	Ln(Hom)	Ln(Hom)	Ln(Hom)
Ln FSS ($t-1$)	0.100** (0.044)	0.107*** (0.037)	0.085* (0.044)	0.086** (0.038)
Ln Rob (t)		0.139*** (0.043)		0.149*** (0.042)
Ln Burg (t)		0.258*** (0.068)		0.226*** (0.072)
Ln percent black (t)			0.233 (0.166)	0.278* (0.164)
Ln %Urb (t)			−0.389** (0.161)	−0.537*** (0.157)
Ln %same house 5 years ago			−10.209*** (0.430)	−0.690 (0.419)
Ln %female headed house			0.790* (0.460)	−0.303 (0.413)
Year fixed effects?	Yes	Yes	Yes	Yes
County fixed effects?	Yes	Yes	Yes	Yes
R^2	0.915	0.921	0.918	0.923
N	3822	3822	3822	3822

Parentheses contain standard errors adjusted for serial correlation (see text). Estimates utilize county population as weight. Analytic sample consists of annual observations for 200 largest counties in U.S. over the period 1980–1999.

* Significantly different from zero at the 10% level.

** Significantly different from zero at the 5% level.

*** Significantly different from zero at the 1% level.

Table 3
Sensitivity analysis

	Ln(Homicide)	Ln(Gun Homicide)	Ln(Non-gun Homicide)
Alternative specifications			
Baseline model, final column, Table 2	0.086** (0.038)	0.173*** (0.049)	−0.033 (0.040)
2 Additional covariates (age, poverty, immigrants)	0.086** (0.036)	0.173*** (0.043)	−0.020 (0.040)
3 Baseline model, unweighted	0.051 (0.043)	0.167*** (0.043)	−0.061 (0.042)
4 Add census division/year fixed effects	0.068* (0.035)	0.162*** (0.044)	−0.047 (0.038)
5 Condition on lag dependent variable	0.061* (0.033)	0.108** (0.046)	−0.032 (0.040)
Alternative samples			
6 Average FSS over 2 years	0.148** (0.059)	0.317*** (0.089)	−0.054 (0.061)
7 Limit sample to largest 100 counties	0.131*** (0.047)	0.207*** (0.066)	0.026 (0.051)
8 Limit sample to largest 50 counties	0.223*** (0.076)	0.252** (0.101)	0.114 (0.078)
9 State-level data, baseline model	0.407*** (0.142)	0.562*** (0.180)	0.106 (0.130)
10 State data, add division/year fixed effects	0.335*** (0.114)	0.534*** (0.167)	−0.066 (0.099)
11 State data, condition on lag dependent variable	0.208** (0.081)	0.272** (0.110)	0.103 (0.116)

Unless otherwise noted, analytic sample consists of 200 largest counties in US using data from 1980 to 1999. Each cell in table presents the coefficient estimate and standard error for the log of FSS ($t-1$) (except for row 6), with the robbery rate, burglary rate, indicators for missing values for robbery and burglary, and percent black as covariates. Parentheses contain standard errors adjusted for serial correlation (see text). Estimates utilize county population (rows 1–8) or state population (rows 9–11) as weights.

* Statistically significant at 10%.

** Statistically significant at 5%.

*** Statistically significant at 1%.

population living in poverty, born outside of the U.S., and in different age groupings – which have almost no effect on the point estimates for FSS. Re-calculating the estimates without weighting by county population produces an elasticity estimate for homicide with respect to guns that is about two-thirds as large as the weighted estimate (row 3). We prefer the weighted estimates because they provide a heteroskedasticity correction. Finally, the results reported in rows 4 and 5 demonstrate that the results hold up quite well to the inclusion of separate year fixed effects for each of the nine Census divisions, or to inclusion of the lagged dependent variable.[8]

The second panel of Table 3 reports the results of several efforts to deal with the fact that our measure of gun prevalence, FSS, is subject to error, primarily due to the relatively small number of suicides in some counties. To increase the “sample” of suicides, we average FSS over 2 years (row 6), limit the analysis to the largest 100 or largest 50 counties (rows 7 and 8), and utilize state-level data (rows 9, 10, and 11). The results suggest that the reduction in measurement error, as expected, tends to increase the point

[8] When we condition on county-specific linear trends (or state trends with the state data), the point estimates for FSS are generally about half as large as in Table 3. These estimates are statistically significant in the state data but not quite significant in the county data, with p-values on the order of $p \approx 0.2$.

Table 4
Specification checks for county and state results, 1980–1999

Outcome	200 Largest county data	State data
Ln(UCR murder)	0.073* (0.043)	0.645*** (0.200)
Ln(UCR rape)	−0.012 (0.048)	−0.201 (0.382)
Ln(UCR aggravated asslt)	−0.040 (0.038)	0.275 (0.168)
Ln(UCR larceny)	0.004 (0.015)	0.096 (0.074)
Ln(UCR MV theft)	0.041 (0.038)	0.046 (0.189)
Ln(Fatality rate from falls)	N/A	0.058 (0.158)
Ln(MV crash fatality rate)	N/A	0.081 (0.068)

Each cell presents the coefficient and standard error (adjusted for serial correlation) for a separate regression of the outcome measure described in the first column against the log of lagged FSS, controlling for the log of the robbery and burglary rates as well as the other covariates described in the final column of Table 4. The county-level regressions condition on county and year fixed effects and weight by county population, using a sample of the 200 largest counties in the U.S.; the state-level regressions condition on year and state fixed effects, as well as weight by state population.

* Statistically significant at 10%.

*** Statistically significant at 1%.

estimates by a factor of from 1.5 to 3 or 4 times our baseline specification. County-level estimates that adjust for measurement error using the approach suggested by Griliches and Hausman (1986), based on a comparison of the within- and first-difference estimators, are also generally about 3 or 4 times those from the baseline model.[9]

A final way to test for the possibility of bias from unmeasured variables is to determine whether FSS predicts outcomes that logically have little relationship to gun prevalence, in the spirit of Altonji et al. (2000, 2002). Table 4 reports the results of estimating the baseline model (final column, Table 2) on rates of other types of crime from the UCR, and on the fatality rate from falls and from motor-vehicle accidents. The estimated coefficients on FSS are not significantly different from zero in any of these regressions.[10,11]

Finally, Table 5 provides suggestive evidence that gun prevalence leads to elevated rates of homicide through the transfer of guns from “legal” to “illegal” owners, rather than through increased gun misuse by otherwise legal owners. In this exercise, we

[9] When we recalculate our estimates with the state-level data using a weighted average of the three gun proxies that are available to us (FSS, gun prevalence from the GSS, and Guns and Ammo subscription rates, where the weights are calculated using factor analysis as in Fryer et al., 2005), the point estimates are about 1.3 times those from our baseline model.

[10] Another implication from Table 4 is that the results are not sensitive to measuring homicides using data from the UCR rather than our preferred source, the Vital Statistics. Note that Duggan (2001) also finds evidence that gun prevalence as proxied by Guns and Ammo subscription rates are not systematically related to other types of crime besides homicide.

[11] We also calculated our estimates using just the long-term variation in gun ownership rates and homicide from the early 80s to the late 90s. This long-difference approach circumvents the problem of modeling the sharp increase and fall of the homicide rate during our sample period. We estimate a long-difference model that shows the changes in log homicides (or log gun or non-gun homicides) from 1980 to 1999, regressed against the change in log FSS over the same period, conditioning on the log changes in the other explanatory variables included in our baseline model. This long-difference estimator yields an elasticity of homicide with respect to gun prevalence of +.3, which is even larger when we pool data from multiple years to correct for measurement error.

Table 5
Effects of gun ownership on Youth Homicides, State Data, 1980–1999

	Ln(Hom 15–19)	Ln(Gun hom 15–19)	Ln(Nongun 15–19)
State data			
Ln(State FSS)	0.593** (0.194)	0.458* (0.205)	−0.053 (0.373)

Each cell presents a coefficient and standard error (adjusted for serial correlation) from a separate regression. Each regression controls for the log of the state's burglary and robbery rate and percent black, log state alcohol consumption per capita, and year and state fixed effects. Estimates are calculated using state populations as weights.

* Statistically significant at 10%.

** Statistically significant at 5%.

focus on homicide rates to victims 15 to 19, a relatively high percentage of whom are killed in gang- and felony-related attacks by youthful criminals—with guns that are typically obtained from the secondary market (Cook and Ludwig, 2004a). That this market is closely tied to the prevalence of gun ownership is suggested by the large coefficient on FSS.[12]

6. Social costs

In sum, gun prevalence is positively associated with overall homicide rates but not systematically related to assault or other types of crime. Together, these results suggest that an increase in gun prevalence causes an *intensification* of criminal violence—a shift toward greater lethality, and hence greater harm to the community. Of course, gun ownership also confers benefits to the owners and possibly other members of the household. The benefits are associated with the various private uses of guns—gun sports, collecting, protection of self and household against people and varmints. But if our estimates are correct, the net external effects appear to be negative.

The magnitude of these net external costs is suggested by the elasticity estimates of homicide with respect to FSS. The baseline model applied to county-level data yields an elasticity of +0.09 or +0.10, although our various attempts to correct for measurement error typically suggest estimates on the order of +0.3 or more. All of these have the feature that the effect on overall homicide is due to changes in gun use, with the possibility of some substitution away from other types of weapon.

These elasticity estimates with respect to FSS also serve as estimated elasticities with respect to the household prevalence of gun ownership, if FSS is proportional to prevalence. Based on cross-section data, FSS does not appear to be strictly proportional—the best-fit line between FSS and survey-based gun ownership rates is linear with a significantly negative intercept (Azrael et al., 2004). But proportionality is

[12] Note that all of the estimates presented here assume that the elasticity of homicide with respect to guns is constant across counties. When we test this assumption by including interactions between FSS and indicators for whether the county's value of FSS in 1980 is in the top or bottom quartile, these interactions are not statistically significant.

a defensible assumption for time-series data: a regression of national handgun prevalence rates (from GSS data) on FSS yields an intercept with a t-statistic of only -1. In what follows, we treat the elasticity with respect to FSS as equal to the elasticity with respect to the prevalence of gun ownership.

The positive elasticity estimates imply that an increase in the prevalence of gun ownership has positive marginal social cost. It is relevant to translate the elasticity into a ratio: the annual change in the homicide count associated with a change in the number of households with guns. That ratio is related to the elasticity by this formula:

$$\text{Ratio of changes in homicides to gun-owning households} = [e \times h \times n]/g \qquad (3)$$

where e = elasticity of homicide rate to prevalence of guns; h = homicide rate per capita; g = household prevalence of gun ownership; n = number of people per household.

This ratio is proportional to the marginal social cost of an additional gun homicide. The formula implies that the marginal social cost of acquiring a gun increases with the homicide rate. For a given homicide rate, the marginal social cost is lower for high-prevalence jurisdictions than low-prevalence—an implication of the log–log specification.

It is important to distinguish between gun types. While handguns make up only about one-third of the private inventory of guns, they account for 80% of all gun homicides and a still-higher percentage of gun robberies. Handguns are also used in most gun suicides. Hence the social costs of handgun ownership are much higher than ownership of rifles and shotguns. Unfortunately, it is difficult to distinguish between the prevalence of long-gun ownership and handgun ownership in aggregate data, since they are very highly correlated across jurisdictions. There is some divergence over time, as overall gun ownership has had a strong downward trend that is not so evident for handgun ownership. FSS is a better proxy over time for handgun ownership.

If the marginal social cost of gun prevalence is entirely attributable to handguns, then the relevant national average is about 20%. Using that value, together with a homicide rate of 10/100,000 (which is close to the average for the 200 counties), an elasticity of +0.10, and 2 people per household, then the formula indicates one additional homicide per year for every 10,000 additional handgun-owning households. In a county with 10% prevalence and a baseline homicide rate of 20, there are 4.0 additional homicides per year for every additional 10,000 handguns; if the baseline homicide rate is 5, and handgun prevalence 30%, just 0.3 homicides are engendered. If the true elasticity is closer to +0.3 instead of +0.1, then the predicted changes in homicides should be tripled.

Two additional questions relevant to calculating marginal social cost cannot be resolved satisfactorily from our results: which margin, and what geographic unit?

6.1. Which margin?

Most households that own one gun own several.[13] FSS is a valid proxy for the prevalence of gun ownership, but much of the "action" is at the intensive margin. With

[13] About three-quarters of all guns are owned by the one-third of gun-owning households that own at least four (Cook and Ludwig, 1996).

respect to providing the right attribution of marginal social cost, it is important to determine whether the acquisition of the nth gun by a gun-owning household has the same cost on average as the acquisition of the first gun. Of course it is only the latter acquisition that will change prevalence.

6.2. What geographic unit?

While our focus has been on county-level ownership, we note that guns often travel across county lines. For that reason, household gun ownership in nearby counties may affect gun availability to local criminals. If true, then "gun prevalence in nearby counties" is a variable that belongs in the homicide regressions, since it is substantively relevant and quite possibly correlated with within-county prevalence. We experimented with specifications that included rest-of-state FSS in addition to the usual within-county FSS, but the results were not very sensible. At this point, it is necessary to be guided by other sorts of evidence regarding the importance of diffuse sources of guns outside of the immediate county. If there are few frictions in the flow of guns to criminals within a state, then our state-level estimates are a better basis for imputing the social costs than the county-level estimates.

Translated into the policy domain, the answers to these questions should influence the nature of regulation adopted in response to the cost argument, and also the geographic scope of the regulatory system. If the number of households with guns, as opposed to the number of guns, is the main concern, then a licensing system may be the preferred form of regulation.[14]

What would be the optimal license fee per household? Answering this question requires monetizing the social costs of the additional homicides that appear to be generated by widespread gun prevalence. One possibility would be to assign each homicide the value per statistical life that has been estimated in previous research, a range of \$3 to \$9 million (Viscusi, 1998), which come primarily from studies of workplace wage-risk tradeoffs. But even the lower end of this range may overstate the dollar value required to compensate the average homicide victim for a relatively higher risk of death, given that (as noted above) such a large proportion of homicide victims are engaged in criminal activity that entails a high risk of death. For example, a study of the wage premium paid to gang members engaged in selling drugs suggests a value per statistical life on the order of \$8000 to \$127,000 (Levitt and Venkatesh, 2000).

Suppose that given local conditions with respect to violence and gun ownership, we estimate a ratio of 10,000 handgun-owning households per annual homicide (approximately what holds at the national average for gun prevalence and homicide with an elasticity of homicide to gun prevalence of $+0.1$) Given a conservative value of life, \$1 million, then the appropriate license fee for a household would be \$100 per year. That license fee would increase with the homicide rate, and in some jurisdictions, such as Washington, DC, would become so high that as to be the practical equivalent of a ban on ownership (a ban on handgun acquisition is currently in place in Washington, Chicago, and some other cities). Of course, this calculation ignores the problem of compliance.

[14] If it is the number of guns that matters, as opposed to the number of households, then an annual tax per gun could be assessed. But our estimates are not directly relevant to estimating the appropriate fee in that case.

This calculation will understate the optimal license fee per gun-owning household if our assumption about the average value per statistical life for homicide victims is too low, or if, as seems likely, gun violence imposes costs on society that are not well captured by any study of the value per statistical life.

Contingent valuation estimates intended to capture the complete social costs of gun violence indicate a value of around \$1 million per assault-related gunshot injury (Cook and Ludwig, 2000; Ludwig and Cook, 2001). On average one in six assault-related gunshot injuries results in death (Cook, 1985; Cook and Ludwig, 2000). Under the assumption that this case-fatality rate is stable across time and space, then at the national averages for gun prevalence and homicide our baseline estimate of a guns/homicide elasticity of +0.10 implies that each additional 10,000 gun-owning households leads to around 6 additional crime-related gunshot injuries. If these contingent valuation estimates are approximately correct, the optimal license fee per gun-owning household would be on the order of \$600. If the true elasticity of homicide with respect to gun prevalence is on the order of +0.30 rather than +0.10, as suggested by some of our estimates that are based on modifications intended to reduce measurement error, the optimal license fee may be as high as \$1800 per household.[15]

Acknowledgements

The research reported here was supported by a grant from the Joyce Foundation, and was conducted in part while the authors were residents at the Rockefeller Foundation's Bellagio Study and Conference Center. Thanks to Mark Duggan for sharing his data on gun magazine subscriptions, to Bob Malme, Dmitri Mirovitski, Joe Peters and Eric Younger for excellent research assistance, and to seminar participants at the NBER Health Spring 2004 meetings, and to David Hemenway, for helpful comments. All opinions and any errors are our own.

References

Altonji, J.G., Elder, T.E., Taber, C.R., 2000. Selection on observed and unobserved variables: assessing the effectiveness of catholic schools. NBER Working Paper, vol. 7831. National Bureau of Economic Research, Cambridge, MA.

Altonji, J.G., Elder, T.E., Taber, C.R., 2002. An evaluation of instrumental variable strategies for estimating the effects of Catholic schools. NBER Working Paper, vol. 9358. National Bureau of Economic Research, Cambridge, MA.

Ayres, I., Levitt, S., 1998. Measuring positive externalities from unobservable victim precaution: an empirical analysis of LoJack. Quarterly Journal of Economics 113 (1), 43–77.

Azrael, D., Cook, P.J., Miller, M., 2004. State and local prevalence of firearms ownership: measurement, structure and trends. Journal of Quantitative Criminology 20 (1), 43–62.

[15] Note that all of these calculations ignore any additional costs that may arise from increased gun prevalence in the form of additional gun accidents or suicides (for estimates of the relationship between gun prevalence and suicide, see Duggan, 2003).

Bertrand, M., Duflo, E., Mullainathan, S., 2004. How much should we trust differences-in-differences? Quarterly Journal of Economics 119 (1), 249–275.
Blumstein, A., 2000. Disaggregating the violence trends. In: Blumstein, A., Wallman, J. (Eds.), The Crime Drop in America. Cambridge University Press, New York, pp. 13–44.
Cook, P.J., 1985. The case of the missing victims: gunshot woundings in the National Crime Victimization Survey. Journal of Quantitative Criminology 1 (1), 91–102.
Cook, P.J., Ludwig, J., 1996. Guns in America: Results of A Comprehensive Survey of Gun Ownership and Use. Police Foundation, Washington, DC.
Cook, P.J., Ludwig, J., 2000. Gun Violence: The Real Costs. Oxford University Press, New York.
Cook, P.J., Ludwig, J., 2004a. Does gun prevalence affect teen gun carrying after all? Criminology 42 (1), 27–54.
Cook, P.J., Ludwig, J., 2004b. The social costs of gun ownership. NBER Working Paper, vol. 10736. National Bureau of Economic Research, Cambridge, MA.
Cook, P.J., Molliconi, S., Cole, T.B., 1995. Regulating gun markets. Journal of Criminal Law and Criminology 86 (1), 59–92.
Davis, J.A., Smith, T.W., 1998. General social surveys, 1972–1998 (machine-readable data file)/Principal Investigator J.A. Davis; Director and Co-Principal Investigator, T.W. Smith; Sponsored by National Science Foundation –NORC ed.—National Opinion Research Center, Chicago (producer); The Roper Center for Public Opinion Research, University of Connecticut, Storrs, CT (distributor).
Duggan, M., 2001. More guns, more crime. Journal of Political Economy 109 (5), 1086–1114.
Duggan, M., 2003. Guns and suicide. In: Ludwig, J., Cook, P.J. (Eds.), Evaluating Gun Policy. Brookings Institution Press, Washington, DC, pp. 41–73.
Fryer, R., Heaton, P.S., Levitt, S.D., Murphy, K.M., 2005. Measuring the impact of crack cocaine. Working paper, Department of Economics, University of Chicago.
Gayer, T., 2004. The fatality risks of sport-utility vehicles, vans and pickup trucks relative to cars. Journal of Risk and Uncertainty 28 (2), 103–133.
Griliches, Z., Hausman, J., 1986. Errors in variables in panel data. Journal of Econometrics 31, 93–118.
Kelly, M., 2000. Inequality and crime. Review of Economics and Statistics 82 (4), 530–539.
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. Journal of Research in Crime and Delinquency 41 (1), 3–36.
Levitt, S., Venkatesh, S., 2000. An economic analysis of a drug-selling gang's finances. Quarterly Journal of Economics 115, 755–789.
Lott, J.R., 2000. More Guns, Less Crime, 2nd ed. University of Chicago Press, Chicago.
Ludwig, J., Cook, P.J., 2001. The benefits of reducing gun violence: evidence from contingent-valuation survey data. Journal of Risk and Uncertainty 22 (3), 207–226.
Solon, G., 1984. Estimating autocorrelations in fixed-effects models. National Bureau of Economic Research Working Paper, No. 32. National Bureau of Economic Research, Cambridge, MA.
Viscusi, W.K., 1998. Rational Risk Policy. Oxford University Press, New York.