Attachment E

Does Strengthening Self-Defense Law Deter Crime or Escalate Violence? Evidence from Expansions to Castle Doctrine

Cheng Cheng†
Texas A&M University

Mark Hoekstra†
Texas A&M University and NBER

Forthcoming in the Journal of Human Resources

Abstract

From 2000 to 2010, more than 20 states passed so-called "castle doctrine" or "stand your ground" laws. These laws expand the legal justification for the use of lethal force in self-defense, thereby lowering the expected cost of using lethal force and increasing the expected cost of committing violent crime. This paper exploits the within-state variation in self-defense law to examine their effect on homicides and violent crime. Results indicate the laws do not deter burglary, robbery, or aggravated assault. In contrast, they lead to a statistically significant 8 percent net increase in the number of reported murders and non-negligent manslaughters.

---

†Cheng Cheng: Texas A&M University, Department of Economics, 3050 Allen Building, 4228 TAMU, College Station, TX 77843 (email: ccheng@econ.tamu.edu). Mark Hoekstra: Texas A&M University, Department of Economics, 3087 Allen Building, 4228 TAMU, College Station, TX 77843 (email: mhoekstra@econmail.tamu.edu).

We would like to thank Scott Cunningham, Steve Puller, Thomas Lemieux, John Winters, Joanna Lahey, Erdal Tekin, Chandler McClellan, Jonathan Meer, seminar participants at the 2012 Annual Meeting of the Southern Economic Association and the 2012 Stata Texas Empirical Micro Conference, and two anonymous referees for providing helpful comments and suggestions. We would like to thank Mark Seaman for providing excellent research assistance.

1. Introduction

A long-standing principle of English common law, from which most U.S. self-defense law is derived, is that one has a "duty to retreat" before using lethal force against an assailant. The exception to this principle is when one is threatened by an intruder in one's own home, as the home is one's "castle". In 2005, Florida became the first in a recent wave of states to pass laws that explicitly extend castle doctrine to places outside the home, and to expand self-defense protections in other ways. Since then, more than 20 states have followed in strengthening their self-defense laws by passing versions of "castle doctrine" or "stand-your-ground" laws. These laws eliminate the duty to retreat from a list of specified places, and frequently also remove civil liability for those acting under the law and establish a presumption of reasonable fear for the individual claiming self-defense. For ease of exposition, we subsequently refer to these laws as castle doctrine laws.

These laws alter incentives in important ways. First, the laws reduce the expected cost of using lethal force. They lower the expected legal costs associated with defending oneself against criminal and civil prosecution, as well as the probability that one is ultimately found criminally or civilly liable for the death or injury inflicted. In addition, the laws increase the expected cost of committing violent crime, as victims are more likely to respond by using lethal force. The passage of these laws may also increase the salience of the legal protections in place, which may itself affect the decision of whether to use lethal force or commit violent crime. The purpose of our paper is to examine empirically whether people respond to these changes, and thus whether the laws lead to an increase in homicide, or to deterrence of crime more generally.

In doing so, our paper also informs a vigorous policy debate over these laws. Proponents argue these statutes provide law-abiding citizens with additional necessary protections from civil and criminal liability. They argue that since the decision to use lethal force is a split-second one that is made under significant stress, the threatened individual should be given additional legal leeway. Critics argue that existing self-defense law is sufficient to protect law-abiding citizens, and extending legal protections will unnecessarily escalate violence. These potential consequences have been of particular interest recently following some highly publicized cases.[1] In examining the empirical consequences of these laws, this study informs the debate over their costs and benefits.

We use state-level crime data from 2000 to 2010 from the FBI Uniform Crime Reports to empirically analyze the effects of castle doctrine laws on two types of outcomes. First, we examine whether these laws deter crimes such as burglary, robbery, and aggravated assault. In doing so, we join a much larger literature on criminal deterrence generally (e.g., Becker, 1968; Ehrlich, 1973; Di Tella and Schargrodsky, 2004; Donohue and Wolfers, 2009). More specifically, however, we join a smaller literature focused on whether unobserved victim precaution can deter crime. For example, Ayres and Levitt (1998) examine whether LoJack reduces overall motor vehicle thefts, while others have examined whether laws that make it easier to carry concealed weapons deter crime (Bronars and Lott, 1998; Dezhbakhsh and Rubin, 1998; Lott and Mustard, 1997; Ludwig,

[1] The most publicized case is that of Trayvon Martin, an unarmed teenager who was shot and killed by a neighborhood watch volunteer (Alvarez, 2012).

1998).[2]

We then examine whether lowering the expected cost of using lethal force results in additional homicides, defined as the sum of murder and non-negligent manslaughter. We also examine the effects of the laws on other outcomes in order to shed light on *why* homicides are affected by the laws.

To distinguish the effect of the laws from confounding factors, we exploit the within-state variation in the adoption of laws to apply a difference-in-differences identification strategy. Intuitively, we compare the within-state *changes* in outcomes of states that adopted laws to the within-state *changes* in non-adopting states over the same time period. Moreover, we primarily identify effects by comparing changes in castle doctrine states to other states in the same region of the country by including region-by-year fixed effects. Thus, the crucial identifying assumption is that in the absence of the castle doctrine laws, adopting states would have experienced changes in crime similar to non-adopting states in the same region of the country.

Our data allow us to test and relax this assumption in several ways. First, graphical evidence and regression results show that the outcomes of the two groups did not diverge in the years prior to adoption. In addition, we show that our findings are robust to the inclusion of time-varying covariates such as demographics, policing, economic conditions, and public assistance, as well as to the inclusion of contemporaneous crime levels unaffected by castle doctrine laws that proxy for general crime trends. This

[2] Our view is that relative to shall-issue concealed carry laws, the potential for castle doctrine law to deter crimes is quite large. For example, in Texas only 1.5 percent of adults age 18 and older have a concealed carry permit, and presumably only a fraction of those carry a gun on a regular basis (Texas Department of Public Safety, 2006; Texas Department of State Health Services, 2006; and authors' calculations). In contrast, Gallup polls indicate that from 2000 to 2009, 44 percent of households own a gun that could be used in self-defense against a burglar or assailant (Saad, 2011). Moreover, strengthened self-defense laws lower the cost of using a concealed carry weapon.

suggests that other known determinants of crime rates were orthogonal to the within-state variation in castle doctrine laws. Along similar lines, we offer placebo tests by showing that castle doctrine laws do not affect crimes that ought not be deterred by the laws, such as vehicle theft and larceny. Failing to find effects provides further evidence that general crime trends were similar in adopting and non-adopting states. Finally, we allow for state-specific linear time trends.

Results indicate that the prospect of facing additional self-defense does not deter crime. Specifically, we find no evidence of deterrence effects on burglary, robbery, or aggravated assault. Moreover, our estimates are sufficiently precise as to rule out meaningful deterrence effects.

In contrast, we find significant evidence that the laws lead to more homicides. Estimates indicate that the laws increase homicides by a statistically significant 8 percent, which translates into an additional 600 homicides per year across states that expanded castle doctrine. The magnitude of this finding is similar to that reported in a recent paper by McClellan and Tekin (2012), who examine these laws' effect on firearm-related homicide using death certificate data from Vital Statistics.[3,4] We further show that this divergence in homicide rates at the time of castle doctrine law enactment is larger than any divergence between the same groups of states at any time in the last 40 years, and that

[3] One advantage of using FBI UCR data is that it allows us to assess both how the laws affect the use of lethal force and whether they deter violent crime. In addition, the nature of the UCR data enables us to measure all homicides, rather than just those caused by firearms. The data also allow us to examine homicide subclassifications and relative changes in reported justifiable homicide from the SHR, along with assumptions about the degree of underreporting, to address the issue of whether the additional homicides are legally justified. The primary disadvantage of the UCR homicide data is that while the annual state-level data we use are regarded as accurate and there is no reason to believe that any total homicide reporting issue at any level should be systematically correlated with changes in castle doctrine law, the monthly data from Vital Statistics are more complete. However, we obtain nearly identical estimates to those reported when we exclude observations in the year in which the state adopted the law, indicating that this is not a problem.

[4] Our findings contrast with those of Lott (2010) in More Guns, Less Crime, who reports that castle doctrine laws adopted from 1977 through 2005 reduced murder rates and violent crime.

3. Identification

To distinguish the effect of the castle doctrine laws from confounding factors, we exploit the within-state variation induced by the fact that 21 states passed such laws between 2000 and 2010. Specifically, we use a difference-in-differences research design that asks whether outcomes change more in states that adopt castle doctrine laws than in states that do not, and focus primarily on within-region comparisons.

Formally, we estimate fixed effects ordinary least squares (OLS) panel data models, where we follow convention and use the log of the outcome per 100,000 population as the dependent variable.[18] For homicide, we also estimate negative binomial models. Ordinary least squares models are estimated with and without weighting by state population.[19] The OLS model estimated is

$$\text{Outcome}_{it} = \beta_1 CDL_{it} + \beta_1 X_{it} + c_i + u_t + \varepsilon_{it}$$

where CDL_{it} is the treatment variable that equals the proportion of year t in which state i has an effective castle doctrine law, X_{it} is the vector of control variables, and c_i and u_t control for state and year fixed effects, respectively. In addition, in most models we also include Census region-by-year fixed effects, to allow states in different regions of the country to follow different trajectories and account for differential shocks by region over time.[20] Note that for states that enacted the law partway through a year, we set CDL equal to the proportion of the year in which the law was in effect, though estimates are almost identical

---

[18] See, for example, Ayres and Levitt (1998), Duggan (2001), and Lott and Mustard (1997). An alternative specification is to use the log of homicide count as the dependent variable, and control for the log of population. Estimates from that specification that correspond to those in column 3 of Table 5 are 0.097 and 0.0602 for weighted and unweighted OLS regressions, compared to estimates reported in Table 5 of 0.0937 and 0.0600.

[19] Specifically, we use analytic weights where average state population over the time period is the weight. This was done using the aweight command in Stata.

[20] There are four Census Regions: West, Midwest, Northeast, and South.

when we exclude the year of adoption.[21] Robust standard errors are clustered at the state level, though we also do additional exercises in the spirit of Bertrand, Duflo, and Mullainathan (2004) to ensure standard errors are being estimated accurately, as well as to perform inference using placebo estimates from pre-castle doctrine data. This last approach of using distributions of placebo estimates to do inference is similar in spirit to the permutation inference approach used in the synthetic control method by Abadie, Diamond, and Hainmueller (2010).

Since we primarily rely on specifications that include state fixed effects and region-by-year fixed effects, the identifying assumption is that in the absence of the castle doctrine laws, adopting states would have experienced changes in crime similar to non-adopting states in the same region of the country. Our data allow us to test and relax this identifying assumption in several ways. First, we look for graphical evidence of whether the two groups diverged prior to treatment. Along similar lines, we offer a formal statistical test by including an indicator in equation (1) for the two years prior to the passage of the laws. That is, we ask whether states that pass the laws diverge even *before* they pass the laws. If they do, it suggests that the identifying assumption of our research design is violated.

We also examine whether time-varying determinants of crime are orthogonal to the within-state variation in castle doctrine laws. Under our identifying assumption, factors such as economic conditions, welfare spending, and policing intensity should not change more over time in adopting states than non-adopting states, as this would suggest that crime

[21] Specifically, when we drop observations containing the year of adoption, estimates corresponding to column 3 of Table 5 are 0.0947, 0.0569, and 0.0895, compared to reported estimates in Table 5 of 0.0937, 0.600, and 0.0879, respectively.

in the two groups might have diverged even in the absence of treatment. Thus, we examine whether adding these controls changes our estimates in a meaningful way. To the extent that our difference-in-differences estimates remain unchanged, it provides some assurance that our research design is reasonable.[22]

Along similar lines, we also show results from specifications that include contemporaneous motor vehicle theft and larceny as controls. While it is possible that castle doctrine laws could affect these crimes, we would expect any such effects to be second-order and at most small in magnitude. Thus, we use these crime measures as controls that pick up any differential trends in crime in adopting and non-adopting states. We also perform falsification exercises using these crimes as outcomes to explicitly test whether castle doctrine laws appear to affect crimes unrelated to self-defense. If our identifying assumption holds, we would expect to see no effects on these crimes.

Finally, we allow for state-specific linear time trends, thereby allowing each state to follow a different trend.

4. Results

4.1 Falsification Tests

One way to test the identifying assumption is to directly examine whether crimes that ought not be affected by the laws—and thus proxy for general crime trends—appear to be affected by the laws.[23] Finding effects on crimes that ought to be exogenous to castle

[22] The primary concern is not that observed determinants vary systematically over time—we can control for those variables directly—but that if they do, it may suggest that unobserved determinants also change systematically over time in the treatment and control groups.
[23] Similar tests are performed by Ayres and Levitt (1998), when they look for effects of Lojack on crimes other than motor vehicle theft.

doctrine law would invalidate our research design.

Thus, we examine whether castle doctrine laws appear to affect larceny or motor vehicle theft. While it is possible that these outcomes are affected directly by self-defense laws, we argue that such effects should be second-order, at best.

Results are shown in Table 3, which uses a format similar to subsequent tables showing other outcomes. Columns 1 through 6 represent OLS estimates that are weighted by population, while Columns 7 through 12 are unweighted OLS estimates. The first column of each group controls for only state and year fixed effects. The second column adds region-by-year fixed effects, while the third column adds time-varying controls. The fourth column additionally includes an indicator variable for the two years before the castle doctrine law was adopted; the fifth drops the leading indicator but adds controls for contemporaneous larceny and motor vehicle theft. Finally, the last column controls for state fixed effects, region-by-year fixed effects, time-varying controls, and state-specific linear time trends.

Estimates for larceny are close to zero and statistically insignificant across all specifications. Estimates of the effect on the log of the motor vehicle theft rate are more interesting. Results in columns 1 and 7 in which only state and year fixed effects are included provide suggestive evidence of increases in motor vehicle theft of 5 to 8 percent, the latter of which is significant at the 10 percent level. However, including region-by-year fixed effects in columns 2 and 8 causes the estimate to drop to zero or even turn negative, and both are statistically insignificant. This suggests that accounting for differences in regional trends in some way may be important in assessing the impact of castle doctrine laws.

4.2 Deterrence

We now examine whether strengthening self-defense law deters crime. We examine three types of crime: burglary, robbery, and aggravated assault. To the extent that criminals respond to the higher actual or perceived risk that victims will use lethal force to protect themselves, we would expect these crimes to decline after the adoption of castle doctrine.

Results are shown in Table 4, where the first 6 columns show estimates from an OLS regression weighted by state population, while the last 6 columns are from unweighted OLS regressions. Results in Column 1 in Panel A for burglary are similar to the finding for motor vehicle theft, in that estimates range from 6 to 8 percent and are statistically significant at the 5 percent level. Again, however, including region-by-year effects in columns 2 and 8 reduces the estimates considerably, and all are statistically indistinguishable from zero at the 5 percent level.

Importantly, there is little evidence of deterrence effects in any specification for any outcome: of the 36 estimates reported, none are negative and statistically significant at the 10 percent level. The estimates are sufficiently precise as to rule out large deterrence effects. For example, in our preferred specification in column 3, the lower bounds of estimates on burglary, robbery, and aggravated assault are -2.1 percent, -1.9 percent, and -2.5 percent. Put differently, our estimates and standard errors from column 3 indicate that if we were to perform this castle doctrine policy experiment many times, we would expect that 90 percent of the time we would find deterrence effects of less than 0.7 percent, 0.4 percent, and 0.5 percent for burglary, robbery, and aggravated assault, respectively. In short, these estimates provide strong evidence against the possibility that castle doctrine laws cause economically meaningful deterrence effects. Thus, while castle doctrine law may well have benefits to those legally justified in protecting themselves in self-defense, there is no evidence that the law provides positive spillovers by deterring crime more generally.[24]

4.3 Homicide

We now turn to whether strengthening self-defense laws increases criminal homicide. Given that the laws reduce the expected costs associated with using violence, economic theory would predict that there would be more of it.

We start by showing the raw data in a set of figures. Figure 1 shows log homicide rates over time for adopting states and non-adopting states, by year of adoption. For example, Figure 1a shows the log homicide rate for the 2005 adopting state (Florida), relative to states that did not adopt the law from 2000 – 2010. While it is somewhat difficult to appreciate the magnitude of changes over time from the graphs and keeping in mind that the adoption year is only partially treated, two patterns emerge. The first is that with the exception of the two states adopting in 2008 (Ohio and West Virginia),[25] the

---

[24] It is worth noting that it is difficult to measure the benefits of these laws to actual victims of violent crime. These benefits could include fewer or less serious physical or psychological injuries, or lower legal costs. We make no attempt to measure these benefits in this paper.

[25] It is little surprise given the small sample sizes involved in this exercise that there would be some set of sets that did not track non-adopting states perfectly in trajectory for the entire period prior to treatment. In addition, we note that while homicide rates did increase in both Ohio and West Virginia from 2000/2001 to 2003, rates there tracked the rest of the country quite closely in changes from 2003 through 2007.

homicide rates of adopting states have a similar trajectory to those of non-adopting states *prior* to the adoption of the law.[26] That is, there is little reason to believe that the homicide rates of adopting states would have increased relative to non-adopting states in the absence of treatment.

Second, Figure 1 shows that there is a large and immediate increase in homicides for states adopting in 2005, 2006, and 2009. Similarly, while the 6 states that adopted in 2007 or 2008 did not appear to experience much of a relative increase in the year of adoption or the year afterward, they notably did not experience the relative drop in homicide rates that other states nationwide did in 2009 and 2010. Of course, given the small samples involved, it is difficult to infer much about short-term versus long-term patterns across these different sets of states, but it is clear from the raw data that castle doctrine states experienced a relative increase in homicides after adoption.[27]

Figure 2 shows the estimated divergence between adopting and non-adopting states over time, where t=0 is the year of treatment. Specifically, Figure 2 graphs coefficients from a difference-in-differences model in which we control for state and region-by-year fixed effects and time-varying covariates, and then allow for divergence 3 and 4 years prior to adoption, 1 and 2 years prior to adoption, the year of adoption, the 1st and 2nd years after adoption, and 3 or more years after adoption. Estimates are relative to

[26] As shown in Figure 1, adopting states have homicide rates that are about 30 percent higher than non-adopting states. However, because we are using a difference-in-differences research design that conditions on year and state fixed effects, differences in *levels* is not a concern for identification. Instead, what would worry us is if the homicide rate in adopting states increased more than in non-adopting states even before treatment, as that would suggest that the groups might have continued to diverge afterward, regardless of the change in law. We see no evidence of that, which suggests that the relative increase seen after 2005 is caused by castle doctrine law. Moreover, note that homicide estimates remained similar even after controlling for time-varying police and incarceration rates and other controls, including region-by-year fixed effects, and allowing for state-specific linear time trends.

[27] We note that estimates remain similar when Florida is excluded from the sample. For example, the estimate from population-weighted least squares declines only slightly from 9.37% to 8.69%, which is still statistically significant at the 1 percent level.

the average difference in log homicide rates 5 or more years prior to law adoption.

Consistent with Figure 1, there is little evidence of divergence in the years prior to adoption. For example, there was almost no divergence in the 4 years prior to adoption using the negative binomial model, and only around 1 to 2 percent using weighted OLS. For weighted OLS, the divergence increases to 10 percent after the year of treatment, and to around 8 percent in the negative binomial model. This offers of preview of the estimated effect on homicide of around 8 percent. There is more modest evidence of divergence prior to adoption using unweighted OLS, though there still appears to be a discrete change at the year of treatment from around 2.5 percent to 7 percent. The difference between the estimated pre-adoption divergence in weighted and unweighted specifications appears to be largely due to the small population states of North and South Dakota.[28]

We now turn to estimating the average effect of the laws in a difference-in-differences regression framework. Results are shown in Panels A, B, and C of Table 5, which show population-weighted OLS estimates, unweighted OLS estimates, and estimates from a negative binomial model. Estimates from the negative binomial regression are interpreted in the same way as those from a log-linear OLS model. Results from the population-weighted OLS model shown in Panel A indicate that the laws increased homicide rates by 8 to 10 percent; all 6 estimates are statistically significant at the 5 percent level, and 3 are significant at the 1 percent level. Estimates from unweighted OLS regressions shown in Panel B range from 5 to 9 percent, though all are

[28] In North Dakota, homicide rates per 100,000 population went from 0.87 in 2000-2002 to 1.5 in 2003-2006, prior to law adoption in 2007. Similarly, homicide rates went from 0.96 in 2000-2001 to 1.89 in 2002 – 2005 in South Dakota, who adopted the law in 2006. South Dakota averages 20 homicides per year and North Dakota averages less than 10, so we suspect the changes in the pre-adoption period were idiosyncratic.

measured imprecisely: t-statistics range from 0.6 to 1.5. Estimates in Panel C from a negative binomial model indicate castle doctrine leads to a 6 to 11 percent increase in homicide. All negative binomial estimates that include region-by-year fixed effects are significant at the 5 percent level, and that which does not (column 1) is significant at the 10 percent level. Finally, we note that homicide estimates are similar for various subsets of the adopting states, as shown in Appendix Table 1. The only difference is the estimates are somewhat larger, albeit not statistically so, for the subset of adopting states identified by Koons (2006) as not previously requiring duty to retreat in either statute or (more typically) case law.[29]

We have also done additional tests in order to ensure that we are making correct inferences about statistical significance. Toward that end, we do tests in the spirit of Bertrand et al. (2004), in which we randomly select 11-year panels from 1960 to 2004, and then randomly assign states to the treatment dates found in our data, without replacement. Thus, we assume that one state expanded castle doctrine on October 1st of the 6th year of the 11-year panel (just as Florida actually adopted in 2005, the 6th year of our panel), and that 13 more states adopted in the 7th year of the 11-year panel, etc. We generate distributions of estimates, and ask how often we reject the null hypothesis of no effect at the 5 percent level, as well as what proportion of the placebo estimates are larger than the actual estimated effect of (real) castle doctrine expansion. The latter figure corresponds to a p-value and is similar to the method used in synthetic control methods (Abadie et al., 2010), as well as by Chetty, Looney, and Kroft (2009).

[29] This is consistent with what one would expect in that states that arguably received a more significant change in law experienced larger (albeit not statistically different) effects. See results in Appendix Table A1.

The resulting placebo distributions from 1,000 random draws are shown in Figure 3, and correspond to Table 5 results from column 2 of Panels A, B, and C, respectively. Results from population-weighted OLS placebo estimates suggest that robust clustered standard errors may be a bit too small: 10.0 percent of simulated estimates are significant at the 5 percent level. However, the estimate of 9.46 percent in column 2 ranks in the 95.4th percentile of placebo estimates, which means only 4.6 percent of placebo estimates are larger than it is.

Results for unweighted OLS simulation results are also interesting. On the one hand, simulations suggest that clustered standard errors from unweighted OLS regressions are accurate: 5.7 percent of the simulated estimates are significant at the 5 percent level. At the same time, however, the estimate of 8.1 percent shown in Table 5 corresponds to the 95.1st percentile, which would give it a p-value of 4.9 percent using the permutation-based approach to inference. This suggests that results in Panel B of Table 5 understate the degree of statistical significance.

Finally, simulations for the fixed effect negative binomial model corresponding to column 2 in Panel C indicate that 7.6 percent of placebo estimates are significant at the 5 percent level, while 14.1 percent are significant at the 10 percent level. As shown in Figure 3, the estimate of 7.3 percent in Table 5 ranks at the 95.7th percentile, as fewer than 5 percent of placebo estimates were larger than the actual estimate in the simulations.

On the basis of these exercises, we conclude that it is unlikely that we would have obtained estimates of the magnitude and statistical significance shown in Panels A, B, and C of Table 5 due to chance.

We have also performed simulations to see if the homicide rates of these particular 21 states *ever* diverged in the way they did after adopting castle doctrine in the late 2000s. To do so, we created 40 panel data sets, each covering separate 11-year time periods between 1960 and 2009. In each 11-year panel, we assume that Florida adopts castle doctrine on October 1st of the 6th year, and that the 13 states that adopted in 2006 adopted in the 7th year, etc. None of the 40 estimates corresponding to either the OLS population-weighted regressions or from the negative binomial regression were larger than those shown in column 2 of Table 5. In the case of the OLS unweighted regressions, only 1 of the 40 placebo estimates was larger than the actual estimate of 8.1 percent shown in Column 2, Panel B, of Table 5.[30] The average estimated divergence across the 40 years was -0.008, -0.004, and -0.005 across the unweighted OLS, weighted OLS, and negative binomial models.[31] Thus, there is no evidence that the homicide rates in castle doctrine states show a general tendency to increase relative to their regional counterparts: in the last 40 years they have almost never done so as much as they did immediately after castle doctrine.

Given the robustness of the estimates to various specifications, it is worth considering what one would have to believe for a confounding factor to cause the observed increase in homicide rates, rather than expansions to castle doctrine. That is, one would have to believe that something else caused homicides to increase relative to non-adopting states immediately after the laws were enacted, but not in the years prior to enactment. In

[30] The one larger estimate was 10.5 percent, and was from the 1975 to 1985 time period.

[31] Estimates for the most recent 5 panels (1995 – 2005 through 1999 – 2009) were 0.022, 0.015, 0.004, -0.027, and -0.069 for weighted OLS, 0.01247, 0.02391, 0.00826, -0.02142, and -0.04719 for unweighted OLS, and 0.004, -0.003, -0.0185, -0.0562, and -0.106 for negative binomial. In these latter panels, we exclude all state-year observations when expanded castle doctrine was actually in effect, so as not to bias placebo estimates upward due to the real treatment effect.

addition, this confounder must have only caused a divergence in homicide rates in the late 2000s coincidental with the passage of castle doctrine law, and not at any point in the 40 years prior. Furthermore, this confounder must cause an increase in homicides in castle doctrine states after adoption, but not cause a similar increase in states *in the same region of the country* that did not expand castle doctrine at that time. Additionally, the confounder must cause adopting states to diverge from their own pre-adoption trend in homicide rate, coincidental with the enactment of castle doctrine law. The confounder must also increase homicides in adopting states after adoption without causing proportionate increases in motor vehicle theft, larceny, robbery, burglary, or aggravated assault. Finally, the confounder must be uncorrelated with changes in the economic conditions, welfare generosity, and the rates of incarceration and policing in adopting states immediately following adoption. We are unable to think of any confounding factor that would fit this description, and thus we interpret the increase in homicides as the causal effect of expanded castle doctrine.

4.4 Homicide: Interpretation

Collectively, we view these findings as compelling evidence that castle doctrine laws increase homicide. However, we note that one downside of the homicide measure is that it could potentially include homicides that are justified under the new self-defense law, but were improperly reported as criminal homicides rather than justifiable homicides. If all the additional homicides were misreported as criminal homicides, the increase may not be viewed by everyone as unambiguously bad. We note, however, that the net increase cannot be driven by a one-to-one substitution of homicides of assailants for homicides of innocent victims. In contrast, in order for the entire increase in homicide to be driven by life-saving use of force, there would have to be at least some cases of multiple killed assailants by a would-be-killed victim.

To shed light on this issue, we look directly for evidence for or against the different interpretations of the increase in reported homicide. We start by examining whether the laws increase the number of homicides classified as murders. This classification available in the Return A files excludes non-negligent manslaughter classifications that one might think would be used more often in potential self-defense killings not classified as justifiable homicides. Estimates in Panel A of Table 6 indicate a similarly sized increase in murder, which suggests that police are largely classifying these additional homicides as murders.

We then turn to assessing whether criminals appear to escalate violence in response to castle doctrine laws. For example, a rational criminal may respond to a real or perceived increase in the likelihood of encountering a victim willing to use lethal force by using a deadly weapon himself. Thus, we examine whether expanded castle doctrine increases felony-type and suspected felony-type murders, which appeared to be committed during a felony. Results are shown in Panel B of Table 6. The estimate from column 1, which controls only for state and year fixed effects, is 10 percent and is statistically indistinguishable from zero. Estimates from specifications including region-by-year fixed effects are more suggestive of a criminal escalation effect: estimates in columns 2 through 5 are around 20 percent and are statistically significant at the 10, 5, 1, and 5 percent levels, respectively, though we note the estimate goes to zero when allowing for state-specific time trends in column 6. We also examine whether criminals are more likely to use guns during robberies.[32] Results in Panel C of Table 6 indicate that there is little evidence of this type of escalation, at least once one compares states to others in their same region.[33] In short, while we find suggestive evidence of escalation by criminals, it is not conclusive.

Finally, we turn to evidence on whether the laws increase the reported number of justifiable homicides. The problem with these data is that justifiable homicides are believed to be underreported: Kleck (1988) estimates that only one-fifth of legally justified homicides by civilians are reported. Only 200 to 300 homicides are classified this way every year in the U.S., compared to around 14,000 total criminal homicides. However, even though the *level* of justifiable homicides may be underreported, *relative* changes in justifiable homicide may still be informative. As a result, we focus on examining the relative increase in reported justifiable homicide, and then estimate how many additional legally justified homicides there really are by scaling the pre-castle doctrine figure by estimates of underreporting.

Results are shown in Panels D and E of Table 6. Panel D shows estimates from unweighted regressions in which the number of justifiable homicides is the dependent variable. Estimated effects range from 1 to 4.3 additional justifiable homicides, which is relative to a baseline average of 4.9 justifiable homicides per state in the year prior to castle

---

[32] We also look at the proportion of assaults in which a gun was used and find no evidence of an increase, though the baseline rate is small (3 percent). We also note that examining these ratios as outcome variables could be problematic if the laws were found to reduce robbery or aggravated assault. However, as we show in Table 4 there is no effect on robberies or aggravated assaults.

[33] It is difficult to think of how using other FBI classifications could help answer this question. For example, the FBI classifies some non-felony-type homicides as having originated in an argument. It is difficult to know, however, whether the argument would have resulted in serious injury to the killer, had that person not used lethal force, or if the argument escalated from, say, a fistfight into a homicide. Yet most would agree that the latter is more disturbing than the former.

doctrine expansion. The estimate in our preferred specification in column 3 is 3.2, is statistically significant at the 5 percent level, and represents a 65 percent increase.[34, 35]

Panel E reports estimates from a negative binomial model. Estimates range from an insignificant 28 percent increase to a significant 57 percent increase.

Using these estimates, we now turn to assessing whether the relative increases observed in Table 6 can explain the entire increase in homicide, given estimates of the degree of underreporting of legally justified homicide. The largest estimated relative increase from a specification in Table 6 that controls for region-by-year fixed effects is 70 percent, which is relative to a baseline total of 103 justifiable homicides across the 21 states in the year prior to castle doctrine enactment. We assume that i) police departments are not *less* likely to report an otherwise-identical homicide as justifiable after castle doctrine expansion, and ii) the relative increase in legally justified homicide due to the change in law is no lower for reporting agencies than for non-reporting agencies. We view the first of these assumptions as likely to hold, and the second as reasonable, though we emphasize that they are in fact assumptions. Combining these assumptions with our estimates in Table 5 suggests that the true castle-doctrine-induced relative increase in legally justified homicide across the 21 states should be no larger than 70 percent.

Kleck (1988) reports that approximately one-fifth of legally justified homicides are reported correctly, while the others are classified as (criminal) homicides. Given the 103 reported pre-castle doctrine justifiable homicides, that suggests that the true figure is 515.

[34] In contrast, we find no evidence of an increase in justifiable homicide by police, consistent with the identifying assumption. Results are shown in Table A2 of the web appendix.

[35] Estimates from weighted OLS are broadly similar. Specifically, estimates corresponding to those in columns 1 through 5 of Table 6 were 9.6***, 6.0**, 4.6*, 4.8, and 4.6*, respectively, where asterisks denote statistical significance. The population-weighted baseline state average was 10.0 justifiable homicides per year.

A 70 percent increase means that castle doctrine expansion causes an additional 361 legally justified homicides, of which 289 (80 percent) would be (mis)reported as homicides. Recall that estimates from Table 5 indicate that castle doctrine law causes approximately an 8 percent increase in homicide, which translates to an additional 611 homicides given the 7,632 pre-castle doctrine homicides. Thus, under these assumptions, our best estimate is that no more than half of the additional homicides caused by castle doctrine law were legally justified.

Of course, different assumptions yield different conclusions. For example, assuming that only 10 percent of legally justified homicides are reported correctly, along with a 70 percent relative increase and the second assumption outlined above, would suggest that all of the additional homicides were legally justified.

To summarize our results, we find no evidence that strengthening self-defense law deters crime. On the other hand, we find that a primary consequence of castle doctrine laws is to increase homicide by a statistically and economically significant 7 to 10 percent. Relative increases in justifiable homicide along an estimate of the degree of underreporting suggest that it is unlikely, but not impossible, that the additional reported criminal homicides consist entirely of legally justified homicides. We emphasize, however, that one's conclusion on that issue depends on assumptions about the nature and degree of underreporting of legally justified homicides.

5. Conclusion

In recent years, more than 20 states have strengthened their self-defense laws by adopting castle doctrine laws. These statutes widen the scope for the justified use of lethal force in self-defense by stating the circumstances under which self-defense is justified and removing the duty to retreat from a list of protected places outside the home. In addition, in many cases they also establish a presumption of reasonable fear and remove civil liability. Thus, these laws could hypothetically deter crime or, alternatively, increase homicide.

Results presented indicate that expansions to castle doctrine do not deter crime. Furthermore, our estimates are sufficiently precise as to rule out moderate-sized deterrence effects. Thus, while our view is that it is *a priori* reasonable to expect that strengthening self-defense law would deter crime, we find this is not the case.

More significantly, results indicate that castle doctrine laws increase total homicides by around 8 percent. Put differently, the laws induce an additional 600 homicides per year across the 21 states in our sample that expanded castle doctrine over this time period. This finding is robust to a wide set of difference-in-differences specifications, including region-by-year fixed effects, state-specific linear time trends, and controls for time-varying factors such as economic conditions, state welfare spending, and policing and incarceration rates. These findings provide evidence that lowering the expected cost of lethal force causes there to be more of it.

A critical question is whether all the additional homicides that were reported as murders or non-negligent manslaughters could have been legally justified. Based on the results of various tests and exercises performed here, our view it is that this is unlikely, albeit not impossible.