The U.S. Consumer Firearms Industry: Welfare and Policy Implications*

Luis Armona†, Adam M. Rosenberg‡

February 27, 2024

Abstract

This paper studies demand for firearms in the United States and its implications for the design of gun policy. We merge the universe of legal firearm transactions in Massachusetts with national price data for individual gun models to estimate a structural model of firearm demand. Our model demonstrates that preferences for firearm characteristics differ substantially along demographic lines, and that consumers are relatively inelastic to price changes. In order to quantify the impacts of firearm policies on public health, we use data on the universe of gun violence incidents to estimate the elasticity of gun violence with respect to gun purchasing. We then extrapolate our demand estimates from Massachusetts to the rest of the United States to predict the impact of alternative gun policies on social welfare. Despite substantial heterogeneity across consumers and firearms, we find that uniform policies on the entire firearms market are still optimal due to substitution at high prices. We estimate that taxes are dwarfed by the benefits of banning firearms purchases entirely. Doubling average prices with an $800 tax on gun purchase would increase welfare by $267 million per year in Massachusetts and $37 billion for the entire U.S. Yet, a total ban on firearm sales would increase welfare by $892 million and $134 billion per year, respectively.

*We are grateful to the Massachusetts Firearms Records Bureau for their assistance in working with the FRB dataset, as well as Mark Bryant and Sharon Williams from the Gun Violence Archive for providing the incident data necessary for this project. We are grateful to Marcella Alsan, Desmond Ang, Jurgen Brauer, David Johnson, Jack McDevitt, and Matthew Miller for their helpful feedback on this paper. Finally, we are also grateful for the financial support via a grant from Arnold Ventures for this research.

†Assistant Professor of Public Policy, Kennedy School of Government, Harvard University. email: larmona@hks.harvard.edu

‡Ph.D. Candidate, Department of Economics, Stanford University. email: arosenbe@stanford.edu

1 Introduction

In the United States, firearms are one of the leading causes of death, and recently became the most common way children die.[1] In 2020, over 45,000 deaths involved the use of a gun.[2] Over 20 million guns are sold each year in the United States,[3] and one in three Americans currently own a firearm.[4] The size of this market and its public health implications makes understanding the structure of the U.S. consumer firearms industry a first-order policy question

Despite its potential policy implications, the literature on firearm markets is relatively sparse. Most studies either analyze existing firearm policies or proxies for gun ownership in relation to their downstream effects on crime and mortality (e.g., Lott and Mustard [1997], Duggan [2001]). The few studies that do analyze market structure tend to use highly aggregated data (e.g., Brauer [2013a], McDougal et al. [2020]). Although these studies provide valuable insight into outcomes of the current firearm market, they provide less information about the substitution patterns of gun consumers, making it challenging to predict market outcomes under alternative regulations.

Relatively few national-level regulations operate on the consumer firearm market. One important policy is determined by the judicial history of the Second Amendment to the U.S. Constitution, which protects the individual right to gun ownership. In 1918, the U.S. Congress passed revenue-raising excise taxes of 10 percent on handguns and long guns,[5] which were recently doubled within the state of California.[6] [7] The National Firearms Act of 1934 and the Federal Assault Weapons bans from 1994-2004 implemented strict controls on the ability of consumers to purchase certain classes of especially dangerous weapons, like machine guns. Although existing research finds that tax policy and weapons bans could dramatically impact the consumer firearms market (Koper and Roth [2001], Moshary et al. [forthcoming]), limited experimentation with national regulation makes it challenging to design efficient firearm policy in the U.S. [Smart, 2018, Smart et al., 2021].[8]

[1] Source: https://www.nytimes.com/interactive/2022/12/14/magazine/gun-violence-children-data-statistics.html

[2] Source: `https://www.pewresearch.org/fact-tank/2022/02/03/what-the-data-says-about-gun-deaths-in-the-u-s/`

[3] Source: `https://www.cnn.com/2021/03/14/us/us-gun-sales-record/index.html`

[4] Source: `https://www.theguardian.com/us-news/2022/dec/15/americans-bought-150m-guns-decade-sandy-hook-shooting.`

[5] Handguns denote firearms such as pistols and revolvers, while long guns consist of rifles and shotguns.

[6] Source: `https://sgp.fas.org/crs/misc/IF12173.pdf`

[7] Source: `https://apnews.com/article/california-guns-ammunition-tax-school-safety-0870a673a3d4e85c78466897cfd7ff6f`

[8] Several other national policies impact the U.S. firearms market. The Federal Firearms Act of 1938 and the Gun Control Act of 1968 jointly impose licensing restrictions on firearm manufactures and retailers, and

This paper studies the structure of the consumer firearms industry in the U.S. and derives implications for the design of gun policy. To understand demand for firearms, we pair administrative microdata on firearm transactions in Massachusetts with 15 years of nationwide data on state-level gun purchase quantities and gun-level prices. We combine these sources of market information to estimate a discrete choice model of legal firearm demand. To study the externalities from firearm purchases, we use an online archive that records incidents of fatalities and injuries from gun violence to estimate the elasticity of gun violence with respect to contemporaneous firearm transactions. Combining our models of firearm demand and gun violence, we evaluate the design of alternative public policies on the U.S. firearms market.

The analysis in this paper is based on a variety of data sources, described in Section 2. We access transaction-level microdata covering every legal handgun and long gun sale in Massachusetts from 2016-2022. Each record includes the buyer's zip code and gender alongside the firearm's manufacturer, model, and product characteristics. Since these data do not record transaction prices, we scrape historical MSRPs from an industry publication, and use these to proxy prices paid by consumers.

In Section 3, we describe features of the data that inform our model of firearm demand. There is considerable variation in the characteristics of firearms available to consumers and meaningful demographic heterogeneity across consumers in different segments of the firearm market. We document that women and individuals in zip codes with larger minority populations are under-represented in the population of potential gun buyers, while individuals in neighborhoods with a conservative voting record are disproportionately likely to buy a high-capacity weapon.

To estimate the price elasticity of firearm demand, we exploit cross-manufacturer variation in firearm production costs. Following McDougal et al. [2020], we use prices of firearm inputs on the world commodity market to construct an instrumental variable based on cost shocks. Our instrument relies on the methodology of Berto Villas-Boas [2007] to isolate cross-manufacturer heterogeneity in exposure to common cost shocks, interacting manufacturer fixed effects with the time series of commodity prices. This results in a high-dimensional manufacturer-by-commodity instrument set, so we apply the dimension-reduction technique of Belloni et al. [2012] to overcome issues of weak identification. Using these instruments, we estimate a price elasticity of aggregate demand for gun purchase in Massachusetts equal to -2.3. This suggests consumers are meaningfully responsive to prices in this market, high-

limit interstate firearm commerce to licensed retailers. In 2005, the Protection of Lawful Commerce in Arms Act extended legal protection to the manufacturers of firearms when their products are used for criminal acts.

lighting the potential of taxes/price shifters as a public policy tool in markets for consumer firearms.

Section 4 uses these reduced-form facts to develop a model of demand for gun purchase. We specify a preference model in the style of Nevo [2000], allowing utility to depend on the interaction of gun characteristics and consumer demographics, price, fixed effects for gun and weapon-class×year, and gun-year demand shocks. We specify the unobservable, idiosyncratic component of utility as a four level nested logit to account for heterogeneous substitution patterns across choices. Our estimated demand model reveals substantial demographic heterogeneity in preferences for firearm characteristics and meaningful differences in preferences for the purchase of any firearm tied to geography, gender, race, and conservative voting patterns. We also find minimal substitutability between handguns and long guns, highlighting distinct motivations for the purchase of each weapon type [Miller et al., 2022].

In order to quantify the public health effects of changes in firearm purchases, Section 5 estimates the causal effect of consumer firearm transactions on downstream incidents of gun violence. To address reverse causality [RAND, 2018], we use the instrumental variable developed by Rosenberg [2023] and isolate variation in firearm transactions due to the net entry of firearm retailers within a zip code over time. We estimate that a 10 percent increase in firearm transactions would lead to a 6.1 percent increase in contemporaneous homicides and 10 percent increases in non-fatal injuries, comparable in magnitude with prior estimates on the effects of gun ownership on gun violence [Duggan, 2001].

To analyze the design of alternative firearm policies, Section 6 uses our estimates of preferences and externalities to simulate the welfare effects of counterfactual changes in gun taxes. In a more conservative set of counterfactuals, we consider policy changes local to Massachusetts. Within Massachusetts, we are able to safely abstract away from industry-wide responses to regulation and to base our predictions on estimates of firearm demand close to our observed transaction microdata. We also pursue a set of counterfactuals based on nationwide changes in gun policy. To extrapolate our estimates of gun demand to the United States, we match our demand model's predictions to aggregate firearm purchases in each state-year via a vertical shifter in willingness-to-pay that varies freely at the state-year level. In these nationwide counterfactuals, we heavily rely on our model to extrapolate away from our data and restrict the potential responses of firearm suppliers in equilibrium. Consequently, we interpret our national-level results more cautiously.

We begin our exploration of counterfactual gun policies by computing the consumer surplus produced by the firearm market under the status quo. Massachusetts is an outlier compared to the rest of the U.S., as its consumers value the firearm market at roughly half of the national average. We find that the average adult in the U.S. derives 121 dollars of consumer surplus annually from the firearm industry. This implies that the firearms industry provides approximately 31 billion dollars of yearly value to U.S. consumers.[9]

We consider a scenario where the regulator has two policy instruments- specific taxes on handguns and long guns, or outright bans. We consider the effects of these levies on the social value of consumer surplus, contemporaneous gun violence, and tax revenue. We find gains to social welfare from increasing taxes on firearms. Low tax rates do little to affect consumers' extensive margin gun purchase decisions, leading to large drops in consumer surplus which are transferred almost one-for-one into gains in tax revenue. At high tax rates, revenue decreases according to the logic of a Laffer curve, but the drop in consumer surplus is more than compensated by a reduction in gun violence. The two tax instruments are complementary in the production of social welfare, as taxing only a single weapon class is meaningfully less efficient. We find that the globally optimal gun policy is a complete ban on all firearm transactions. A gun ban would destroy all 31 billion dollars of value that the gun market provides to consumers each year nationally. However, a complete ban on purchases would save society more than 160 billion dollars through the reduction of gun violence, netting a 134 billion dollar increase in social welfare. Although a ban on firearm transactions is unlikely under the current U.S. political climate [Gentzkow et al., 2019, Luca et al., 2020] and constitutional law, we view this result as demonstrating the social welfare gained by other industrialized countries that design restrictive policies on their firearm markets. Our health predictions coincide with the observed effects of Australia's 1996 National Firearms Agreement, which implemented broad-scale gun bans across the country [Ramchand and Saunders, 2021]. Moreover, our preference estimates suggest that Australia's sweeping gun policy likely increased social welfare.

Most of the existing literature surrounding firearms has focused on estimating the relationship between guns and crime or violence Duggan [2001], Cook and Ludwig [2006], Koenig et al. [2016], Johnson and Robinson [2021], Studdert et al. [2022]. The evidence is somewhat mixed on how guns relate to public health, but these studies generally point to a positive relationship between gun ownership and violence [RAND, 2018]. We contribute to this literature by utilizing our high resolution data on both legal firearms transactions and gun violence incidents to estimate the effect of firearm transactions on fatal and non-fatal injuries from gun violence. We also combine these estimates with our demand model to study the design of firearm regulations across different weapon classes at the state and national level.

This paper most closely contributes to a growing body of work that analyzes the design

[9] Comparable estimates from Grieco et al. [2023] value the new car market in 2016 as producing 150 billion dollars of consumer surplus.

of regulation on consumer firearms markets in the U.S. [Koper and Roth, 2001, Bice and Hemley, 2002, Knight, 2013, McDougal et al.]. We advance the literature on the consumer firearms market by using revealed preferences to estimate a micro-founded model of demand for guns in characteristic space and by using this model to make welfare comparisons against policy alternatives.

Our work is especially close to two recent papers within the literature on gun market regulations. Moshary et al. [forthcoming] use a conjoint experiment and stated preferences to estimate a micro-founded model of demand for guns and examine the implications of national firearm policies on consumer surplus. They estimate a low price elasticity for these consumers, consistent with our findings. Rosenberg [2023] assembles 40 years of firearm transaction records from California to estimates a nested logit model of gun retailer preferences and externalities from firearm ownership, in order to evaluate counterfactual regulations on local firearm markets. We view our paper as applying the analysis of high-resolution data on firearm markets and gun violence from Rosenberg [2023] to the questions of national regulation posed by Moshary et al. [forthcoming]. We are also the first paper to our knowledge to use firearm transactions microdata to estimate the price elasticity of gun demand.

2 Data

We now describe the data elements used in the paper.

2.1 FRB Firearms Transactions

Our primary dataset is the Massachusetts Firearm Record Bureau's (FRB) repository of firearm transaction records. For each purchase from a Federal Firearms Licensed (FFL) dealer, the state of Massachusetts requires an electronic verification that the consumer has the appropriate license to purchase the firearm. Thus, this dataset captures the universe of legal gun transactions occurring in the state of Massachussetts.[10]

Before completing a transaction, dealers are required to record information about the firearm, which are manually input as text by the dealer. These include the make (manufacturer/brand), model of the gun, type of weapon (handgun, rifle, or shotgun), the caliber of the gun, the barrel length, the surface finish of the weapon, and whether it is a high capacity weapon. Caliber denotes the width of the barrel of the gun, and approximately measures the power of the weapon. Barrel length captures the range and precision of the firearm.

[10] We use a publicly available version of this data, available here: `https://www.mass.gov/info-details/data-about-firearms-licensing-and-transactions`

High capacity designates weapons capable of accepting high capacity ammunition devices. Surface finish denotes the material/color on the exterior of the gun. All of this is available in the FRB transactions data. The data includes information on the dealer/seller of the firearm, such as their FFL number, the zip code of the retail store's location, and the name of the store. It also includes data on the buyer, such as their gender, and the zip code they reside in according to their firearm license. Finally, the transaction date is also recorded. the FRB began recording these transactions in 2006, and began tracking whether weapons are high capacity in 2016. As such, we use data on transactions from 2016-2022. There are 816,000 transactions from retailers recorded in this time period.

To process this data, we convert calibers to a standard unit of inches across all gun models, as caliber is recorded as inches, millimeters, and gauge for shotguns. We do a similar exercise for barrel length. We manually standardize the manufacturer field, for any manufacturers with at least 30 transactions in the data. We limit most of our analysis to the top 100 makes; these consist of 92% of all transactions in the data.

2.2 Blue Book of Gun Values

Our second dataset is a panel dataset of gun models and their historical prices from the Blue Book of Gun Values (BBGV). Through purchasing a subscription, we were able to collect all available information on guns in the BBGV.[11] For each gun model, BBGV has annual historic prices from 2006-2022, with pricing information varying by the condition of the firearm. For new guns, the BBGV lists the MSRP (manufacturer's suggested retail price), and indicates the gun is currently in production. Used gun prices are distinguished by the grading (condition) of each model, ranging from 10% to 100%. Moshary et al. [forthcoming] notes that there is little geographic variation in gun prices, so MSRP may be a good indicator for the national prices of guns across local markets. We use this as our primary measure of price in this paper. BBGV is organized hierarchically by gun manufacturer, then gun type (e.g. semi-automatic pistols), and then gun model. We manually convert BBGV weapon types to the three types (handgun, rifle, and shotgun) that appear in the FRB. It also includes text metadata on gun models and gun types within each manufacturer. In total, BBGV contains price data on 9,962 unique gun models, across 71 manufacturers. Price data is somewhat incomplete across gun models: MSRP is available for only 27% of gun-year observations, mostly because many guns in the database are no longer manufacturered, and primarily purchased as collectors' items.

Because the gun model names in the FRB dataset are not standardized, we merge this

[11] Downloaded from `https://www.bluebookofgunvalues.com`

price dataset to the FRB transactions data using a fuzzy string matching algorithm approach. We perform this merge using a multi-step approach, which is described in detail in Appendix A. In total, through this procedure, we are able to match 90% of FRB transactions to a BBGV ID, for a total of 7,616 unique BBGV gun models. About 70% of guns are purchased in years which they are actively produced. We assume for the remainder of the paper as if these actively produced guns were purchased new from a retailer and sold at the MSRP price.

2.3 Gun Violence Archive

We measure the downstream impacts of consumer firearm ownership using proprietary microdata from the Gun Violence Archive ("GVA").[12] The GVA is an independent, online repository of incident- and individual-level data on gun violence events. It covers fatal and non-fatal incidents of assault, accidental, and officer-involved gun violence but is not designed to record most self-harm incidents.

Data in the GVA are gathered each day from a combination of automated queries, manual readings of local and state police records, and other sources including media and government information. For inclusion in the GVA, an incident must be verified by an initial researcher and validated through an addition process. Data in the GVA is not crowdsourced or recorded via bots. Although the data collection process has changed over time, discussions with GVA staff suggest that the collection process has been roughly constant since 2014.

in June 2023, we obtained an extract of all records in the GVA.[13] Figure A1 benchmarks counts of gun violence incidents in the GVA against publicly available records of other-directed gun fatalities from the CDC. Beginning in 2014, we find an exceptionally close match between counts of gun-related assault fatalities in these two datasets. The CDC does not record non-fatal firearm injuries, but best estimates suggest two non-fatal incidents occur for every fatality (Kaufman et al. 2021), a ratio which we replicate in the GVA. We thus treat the GVA as representing the universe of gun violence incidents that are not self-inflicted, whether fatal or not.

Each recorded incident in the GVA includes its associated counts of fatalities and injuries, the incident date, and the geocoded address at which the incident occurred. We aggregate this microdata to compute counts of gun violence incidents, non-fatal injuries, and fatalities at the zip code-year level.

[12] `https://www.gunviolencearchive.org/`

[13] We gratefully thank Mark Bryant and Sharon Williams for access to this data.

2.4 Auxiliary Datasets

We complement these main datasets with auxiliary data. In order to understand demographic heterogeneity of buyers, we merge demographic data of buyer zip codes from the 2019 5-year American Community Survey (ACS).[14]. This includes the distribution of gender, age, race, education, poverty in the zip code, along with the median zip code income and the density (individuals per square mile).

Given the political nature of gun ownership [Joslyn et al., 2017], we also include a measure of local political ideology. For this, we download the precinct-level vote shares in the 2016 U.S. presidential election, from the Voting and Election Science Team [Voting and Election Science Team, 2018].[15] We aggregate precinct-level vote shares for all conservative candidates in the presidential election[16] in 2016 to construct a "Percent Conservative" measure for each precinct in the United States. We then aggregate this to the zip-code level using the supplied VEST precinct shapefiles along with U.S. Census TIGER ZCTA shapefiles, weighting precincts by total number of votes.

Though our data on gun transactions comes only from the state of Massachusetts, we aim to understand national implications of gun policy. to extrapolate our data to other states, we use data on the number of guns sold in each state-year from the FBI's National Instant Criminal Background Check System[17]. For each legal firearm purchase in the United States, a criminal background is required. This is adjusted using the methodology of Brauer [2013b] to produce an estimate of the total number of gun sales in each state-year.

Not all individuals in a geographic area are relevant to the market for firearms. Given that firearms are a polemic public health topic in the United States [Oliver, 2006], some individuals may be ideologically opposed to owning a firearm, and moreover, this likely varies from region to region of the United States, as well as across gender. To account for this, we utilize wave 26 of the American Trends Panel (ATP), conducted by Pew Research, to better identify the set of potential gun purchasers in each market. Specifically, we utilize a question that asked 3,900 respondents whether they owned a gun, or could see themselves owning a gun, as an indicator of being in the market for a gun. We map these to the observed demographic characteristics of the respondents via a LASSO regression [Tibshirani, 1996], to create an estimate of the potential market size in each zipcode-gender cell. This is similar to the approach in Backus et al. [2021] to predict market size for cereal at local grocery stores.

[14]Downloaded from the IPUMS NHGIS page: `https://www.nhgis.org/`

[15]Downloaded from `https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/NH5S2I`

[16]This includes Donald Trump of the Republican Party, Gary Johnson of the Libertarian Party, Evan McMullin (Independent) and Darrell Castle of the Constitution Party

[17]Downloaded from `https://github.com/BuzzFeedNews/nics-firearm-background-checks`

Explicitly, we regress (by census region) the indicator for potential gun buyer on the full set of Pew-provided indicators for gender, age category, education level, race/ethnicity, citizenship, marital status, income category, health insurance, and party affiliation. The LASSO then selects the non-zero variables to include in the linear probability model. We then project these coefficients onto the observed zip code demographics to obtain our estimates of market size, which are about 1/3 of the adult population in each zipcode. Estimates are shown in Appendix Figure A3. Male, U.S. Citizenship, and Republican party affiliation are the largest predictors of willing to own a gun. Moreover, the importance of demographics vary across different regions of the United States.

Our goal is to use the rich transaction-level data in Massachusetts to learn about demand in the United States. To this end, we also use the NICS background check data as our measure of demand for guns in other states in the United States.[18] These background checks occur each time a consumer attempts to buy a firearm from a nationally licensed gun dealer (FFLs), though in some states, gun permit holders can present their permit in place of the background check, so coverage varies from state to state [Lang, 2016]. At the same time, the FBI reports these data at a state-year level, and in terms of national coverage, is the highest fidelity data researchers have access to, so it is a commonly used proxy for the flow of guns from suppliers to consumers. In Appendix Figure A2, we compare the sales of firearms recorded in FRB to the number of background checks, and show, at least in Massachusetts, that they are very similar.

2.5 Summary Statistics

Table 1 displays summary statistics on our set of matched gun models from BBGV. Handguns are typically shorter in length, and cheaper on average than long guns (rifles and shotguns). On average, prices for a new gun are around \$1000, though there is significant price variation even within a class of weapons, as indicated by the standard deviation upward of \$600. About 40% of long gun models are shotguns. Figure 6 plots the distribution of barrel length and caliber within long guns and handguns. Each point represents a gun model, with size proportional to its (logged) purchase frequency in the FRB data. While historically, analysis of this market has aggregated to the type of gun (handguns, long guns), there are appear to be significant variation in the characteristics, even within a weapon class. The extent to which these characteristics may matter for consumer choice in this market is an empirical question we will visit in our demand model.

Figure 1 plots the median and interquartile range of new gun prices (MSRP) for our

[18] Data downloaded from `https://www.fbi.gov/file-repository/nics_firearm_checks_-_month_year_by_state.pdf/view`

set of gun models over time. The plot highlights a great deal of variation in gun prices, suggesting there is a role for product differentiation within these weapon types, since some gun manufacturers are presumably charging different prices to cater to different consumers. Consistent with Table 1, there is greater price variation in the long gun category. Prices for guns have been increasing over time, and this is not simply due to different sorts of guns being offered by manufacturers, or the variety/mix of guns offered in a particular. Figure 2 shows the estimated year fixed effects of a two-way fixed effects regression on log(MSRP), partialling out the time-invariant price levels of each gun model. Estimates are shown relative to 2016, the base year. For both types of weapons, gun prices remained relatively stable until the start of the pandemic, at which point prices increased substantially: within model, in 2022, prices were around average 8% higher than 2 years prior. This is possibly attributable to the well known supply chain issues that plagued many industries in the United States, and in particular, the firearms manufacturering industry, which devotes a large share of its expenditures to purchasing of inputs such as steel and copper[19]

In Table 2, we display the demographic characteristics of gun buyers in this market, based on their gender and the zip code they reside in. We compare this to the demographics of the Massachusetts adult population in Column (1). In Columns (2), we filter to those predicted to be "in the market" for potentially purchasing a firearm, according to our linear probability model based on the ATP questionnaire. In columns (3)-(5), we weight each demographic based on the number of purchases of guns with the associated characteristic. Relative to the adult population, gun buyers are more likely to be male, and live in a less dense and richer zipcode, as well as zipcodes with a lower fraction of racial minorities. These demographic patterns are more stark for long gun buyers relative to handgun buyers.

In Figure 3, we plot the aggregate number of purchases over time in Massachusetts. About 70,000 handguns 40,000 long guns are purchased each year. Gun sales appear to be declining over time during our sample, until the COVID pandemic, at which point there is a spike in purchases that coincides with the observed price increases in guns. If we split this by capacity of the weapon (Figure 4), about 30,000 high-capacity guns are sold each year across handguns and rifles.[20] In Figure 5, we plot new versus used gun purchases (defined as those with an MSRP in BBGV for a particular year). We see that the new guns are more popular, and while both types of guns experience an increase in sales after the COVID-19 pandemic, used guns in particular experience a greater surge.

[19]See, for example, this article: `https://shootingindustry.com/discover/supply-chain-woes/`

[20]High capacity shotguns are extremely rare.

3 Descriptive Evidence

3.1 Heterogeneous Demand for Firearms

We begin our analysis by exploring the demographic correlates associated with firearm demand, to uncover whether consumers of firearms have observably heterogeneous purchasing patterns.

First, we correlate gun purchases frequencies with consumers demographics. To do so, we collapse the transaction data to the zip-gender-month (z, g, t) level, our most granular measure of consumer demographics. If some gender-female demographic cells contain no data for a particular month, we infer zero purchases for that period, since the FRB data contains all legal transactions in the state. We then estimate a regression of the following form:

$$\frac{\text{Purchases}_{z,g,t}}{\text{Pop}_{z,g} \cdot 1000} = \beta D_{z,g} + \alpha_{t,\cdot \#\text{Dealers}_{z,t}} + \epsilon_{z,g,t} \tag{1}$$

where $D_{z,g}$ denotes cell-level demographics, and the LHS variable is the number of purchases per 1,000 adults in each month. We include in this regression a set of $year - month$ fixed effects, to partial out temporal patterns in gun purchases, interacted with a fixed effect for the number of licensed gun retailers[21] in the zip code that same month. These latter set of controls serve to better control for the "market structure" that consumers face during any particular purchase decision, in the style of Bresnahan and Reiss [1991]. This is done to isolate the potential variety/choice set consumers may face in their local retail setting. Thus, the thought experiment is as follows: Holding the variety of guns available constant, how does the purchase frequency of consumers with demographics $D_{z,g}$ differ in their likelihood from the population average? This serves to capture heterogeneity in the "extensive margin" of gun purchases.

We include in $D_{z,g}$ a dummy variable for being female, the fraction of conservative voters as of 2016, the fraction white, the fraction with at least a bachelor's degree, as well as the (logged) median income and population density. The gender variable comes from the FRB, the political affiliation measure comes from the VEST dataset, while the rest of the demographics come from the 2019 5-year ACS.

Next, we analyze the demographic heterogeneity in the characteristics of weapons that consumers choose, conditional on a purchase. To do so, we utilize the tranasaction-level

[21] Measured as the number of licensed gun retailers stores open in that zip code, according to the FRB data.

micro data, and estimate an analogous regression:

$$X_{i,j,t} = \beta D_i + \alpha_{t,\#\text{Dealers}_{z,t}} + \epsilon_{i,t} \tag{2}$$

Where $X_{i,j,t}$ denotes a characteristic of the gun consumer i chose to purchase, and α denotes the same fixed effects as Equation 1. Note that because this directly utilizes the transaction-level data, this regression is done conditional on purchasing a weapon, so differs qualitatively from the previous exercise. The corresponding thought experiment to Equation 2 is: conditional on purchasing a gun, holding the assortment of guns available constant, how do the attributes of the gun a consumer with demographics D_i purchases differ from the average gun characteristic in the data? Correspondingly, we label this the "intensive" margin of consumer heterogeneity in gun purchasing.[22]

Figure 7 displays the estimates from these regressions. Standard errors are clustered at the zip-code level. In Panel (a), we plot estimates of β in Equation 1. Note that we scale the effects of logarithmic variables so they are of similar magnitude to effect of binary/fraction demographics. The panel shows that areas with more white adults are correlated with more gun purchases, controlling for market structure, suggesting that these consumers enjoy purchasing firearms more frequently. The coefficient on female is precise and negative, consistent with the findings from the ATP Survey that women are less likely to be willing to own a gun.

Panel (b) of the Figure presents demographic correlates with gun characteristics. On the left side, we display correlates with the discrete type of weapon (handgun, rifle, shotgun, or high capacity), while the right side shows correlates with caliber and barrel length of the weapon (in terms of the sample standard deviations of these attributes). Conditional on purchasing a firearm, women are more likely to purchase handguns, while those in white neighborhoods are more likely to purchase traditional low-capacity long guns. We also find that those in conservative neighborhoods are more likely to purchase high capacity weapons. Consumers living in more urban and affluent areas are more likely to buy smaller weapons.

We note that this analysis is purely descriptive in nature, even though we include a rich set of fixed effects to better isolate demand-side responses. Nonetheless, the results show meaningful demographic heterogeneity in the types of weapons firearms consumers choose to purchase, and are highly suggestive that demographic heterogeneity may play a significant role in firearms preferences, beyond propensities to purchase guns. We incorporate this demographic heterogeneity into our demand-side model of preferences in Section 4.

[22]Typically, intensive vs extensive margin refers to choosing to buy, versus how many products to purchase. Because we lack individual identifiers in our data, we cannot perform this exercise; so Equation 1 will capture both extensive and intensive margins under that more standard terminology.

3.2 Demand Price Response

We now describe the impact of price on consumer firearm demand. This is a critical object for the design of firearm policy, especially taxes, but challenging to estimate because price changes in the data may be correlated with unobservable (to the econometrician) demand shocks. Such simultaneity generates price endogeneity, requiring the usage of instruments that shift price but are uncorrelated with demand.

We use input cost shocks as instruments for price, a popular approach in the demand estimation literature [Berry and Haile, 2021]. This instrument is likely to work well in the firearms industry, where much of the costs are tied up in the purchase of production materials, such as steel and aluminum. According to the 2017 BEA Input-Output Tables, intermediate inputs account for 63% of costs, and among intermediate inputs, 35% of purchases are on primary metals (e.g. steel) and fabricated metal products (e.g. iron molds).[23]

Because we do not observe the exact mix of material used in firearm production, we propose a strategy to construct cost shocks in the spirit of Berto Villas-Boas [2007], by flexibly approximating manufacturer-specific production functions. Explicitly, our approximation endows each manufacturer with a unique response to a set of potential metal inputs by interacting manufacturer IDs with a price index for each metal. We use these interactions to predict firearm prices, and then use the predicted prices as instruments for the observed prices in the sales data. This mirrors the strategy of McDougal et al. [2020], which uses the price of hot and cold rolled steel to estimate the price response of aggregate firearm demand.

For our set of potential inputs, we use the prices of primary metal commodities from the Bureau of Labor Statistics' monthly PPI commodity data series.[24] We aggregate these to annual price indices with 2016, the first year of our sample, as the base year. This provides us with 51 time series for the prices of potential firearms inputs, ranging from carbon steel scrap to copper ore. Our rationale for using these commodity prices as instruments is that these raw materials are traded in global commodity markets, and so changes to the price of these inputs are unlikely to be related to the demand for firearms in Massachusetts. Our

[23] Source: author's calculations using intermediate inputs for "Ammunition, arms, ordnance, and accessories manufacturing" commodities (excluding inputs from the same commodity type). Data downloaded from `https://www.bea.gov/data/industries/input-output-accounts-data`

[24] See `https://download.bls.gov/pub/time.series/wp/wp.txt` for a description of the dataset. We use all 6-digit commodity series beginning with 101 and 102, denoting iron/steel and nonferrous metals, respectively.

empirical specification is as follows:

$$\log(p_{j,t}) = \gamma_{j,1} + \eta_{c,t,1} + \sum_{k=1}^{51}(\beta_{m(j),k}\log(\tilde{p}_{k,t})) + \epsilon_{j,t} \tag{3}$$

$$IHS(q_{j,t}) = \gamma_{j,2} + \eta_{c,t,2} + \alpha\log(p_{j,t}) + \xi_{j,t} \tag{4}$$

where $p_{j,t}$ is price, γ_j denotes gun model fixed effects, $\eta_{c,t}$ denotes year $\times$ weapon class (long gun versus handgun) fixed effects, $\tilde{p}_{k,t}$ denotes the price index in year t of input k, $\beta_{m(j)}$ denotes a coefficient on input prices allowed to vary by manufacturer m of gun model j, and $\xi_{j,t}$ is the unexplained component of sales, stemming from, for example, unobserved demand shocks. To take advantage of the full sample of guns with purchase data in Massachusetts, we use the inverse hyperbolic sine in place of the log transformation, to retain those guns we know are available (listed MSRP that year), but report zero sales in MA (about 1/3 of our gun-year observations).[25] Assuming the log approximation is accurate, α represents a price elasticity, and captures both extensive (buying no gun) and intensive (switching to a different gun model) margins of demand.

The specification in Equations 3 and 4 may be estimated using a two-stage least squares approach, instrumenting for the price of each gun with the manufacturer-specific production function. Besides ensuring relevance, for the price coefficient to be correctly identified, we must also assume that the excluded instruments only affect the demand through their impact on price.

Figure 8 shows the annual variation in the 51 input price indices used for identification, grouped by the 2-digit code of each commodity. As we can see, there is large variation in the evolution of these commodity prices, with many commodities doubling in price since 2016. The largest shock to these input prices coincides with the COVID-19 pandemic, which spurred a series of supply chain issues across the globe. COVID also coincides with a large increase in gun sales [Sokol et al., 2021]. In our estimation, we do not want to ascribe potential demand shocks due to COVID to changes input prices. For this reason, our year fixed effects η are critical to identification.

With year fixed effects in place, the residual variation captured by our instruments is the *heterogeneous exposure* of manufacturers to input price shocks. For example, if a manufacturer primarily uses carbon steel scrap to produce firearms, they will be more exposed to the jump in the carbon steel scrap price in 2021, relative to a manufacturer that uses other metals during production, and so would have a positive coefficient $\beta_{m,\text{carbon steel scrap}}$. The primary threat to identification would be if demand shocks, such as COVID-19, were

[25] Results are very similar if we use the $\log(x+1)$ transformation.

correlated with the reliance of certain manufacturers on particular inputs. For example, if the surge in gun demand during COVID-19 was driven by a surge in demand for guns made out of carbon steel scrap specifically, then our estimation procedure would violate the exclusion restriction. As additional protection against this type of simultaneity, we note that firearm pricing is likely to respond to national demand shocks, while we estimate demand from consumers in Massachusetts, which accounts for only 0.8% of national firearms sales.[26] Thus, even if our exclusion restriction were to be invalid at the national level, it may still hold within Massachusetts, where demand shocks are unlikely to affect pricing strategic.

We restrict the estimation sample to those 850 models that have at least 10 purchases during our sample period; this is about 85% of all transactions on new guns, and represents 99 manufacturers in the data. This is the same restriction we make when estimating our demand model in Section 4. As a result, our set of instruments is over 5,000, more than the number of gun-year observations in our sample, and consequently our first-stage would not be identified. Even if we pruned the set of commodities to have a manageable set of instruments, we would worry about weak instruments impacting our identification. For some manufacturers, production processes may be standardized and mimic those of their competitors, yielding little useful variation in production.

For this reason, we apply the routine described in Belloni et al. [2012] and estimate the first stage via a Post-OLS LASSO IV. The LASSO is a widely used machine learning algorithm that implements a L1 penalty on regression coefficients [Tibshirani, 1996], both downweighting the magnitude of coefficients, and allowing for some coefficients to be set to zero. Under certain conditions, this algorithm has an "oracle" property, meaning that it can recover the true set of relevant covariates among a large set of candidate covariates. The procedure in Belloni et al. [2012] uses a data-driven penalty based on econometric theory that, when used as the first stage to select the instruments, can be used in a standard 2SLS IV approach and the conventional standard errors are valid. Explicitly, the procedure is as follows:

1. Using the data-driven penalty to recover the set of non-zero instruments, allowing for independent but heteroskedastic errors.

2. Run a post-LASSO OLS of price on the selected instruments.

3. Construct $\hat{p}_{j,t}$ as the fitted values from the post-lasso OLS regression.

4. Estimate the second stage using $\hat{p}$ as the instrument for the observed prices.

[26]Calculated from NICS Background check data.

Throughout, both the LASSO and IV are run with time and gun model fixed effects always partialed out.

Figure 9 displays the selected (non-zero) coefficients from our baseline specification with gun model and weapon-class times year fixed effects (Panel A), along with the fit of the annual change in LASSO predicted prices versus the annual change in observed MSRP (Panel B), indicating a good fit. From the potential set of 5049 instruments, 17 are selected, suggesting that the pruning done in the first stage is important and many of the manufacturer-specific coefficients are uninformative. Unsurprisingly, the 17 instruments selected in the first stage are associated with the largest manufacturers in the data. The selected instruments shift the prices of 40% of the gun models in our sample, ensuring our elasticity estimates are not driven by the price changes of only a few models.

Table 3 shows the results for the Post-LASSO IV in the reduced form. In Column (1), we report the OLS estimates from regressing purchases on price using two-way fixed effects. Perhaps unsurprisingly given the insignificance of Massachusetts in the national market for firearms, we estimate a negative price elasticity, but it is statistically insignificant and small in value (-0.59). In Column (2), we report the Post-LASSO IV estimate. Here, the point estimate is larger in magnitude. The estimated price elasticity changes qualitatively, to -2.3, and is now statistically significant at conventional levels. Because the standard F-test for relevance is invalid, due to the fact that we are estimating the IV using those instruments that were found *ex-post* to be important, we instead use the recommended sup-score test statistic from Belloni et al. [2012], which tests for joint significance of the excluded instruments, accounting for LASSO selection. The test confirms relevance as the estimated test statistic of 14 is well above the critical value of 4.56 for 1% significance. Column (3) adds class $\times$ year in place of year fixed effects, allowing for differential time trends in demand for long guns and hand guns. The point estimate is similar to Column (2).

The LASSO-selected IV specifications suggest a price elasticity for firearms of around -2. This estimate is more elastic than the estimates of Moshary et al. [forthcoming], but similarly small in magnitude. These estimates suggest that demand is somewhat price inelastic. Consequently, gun taxes acting as price shifters may be a less effective policy tool for changing behavior in this market. Moreover, we are unable to separate differential types of substitution patterns (exiting the market versus buying a different weapon) in our reduced-form setup. With this limitation in mind, we now introduce a structural model of demand, that imposes stricter assumptions on consumer behavior, but in turn allows us to better isolate these channels, and evaluate the efficacy of certain firearm regulations.

4 A Model of Firearm Demand in Massachusetts

We estimate a discrete choice model in the style of Berry et al. [1995] based on transaction data in Massachusetts. The model accounts for heterogeneous preferences based on gender and characteristics of the neighborhood consumers reside in, which is the full set of demographic data available from the FRB. We also allow for unobserved differences in individual-preferences in sub-groups of products, as well as idiosyncratic errors, via a nested logit specification. We then extrapolate these parameters to the rest of the United States using moment matching conditions, assuming consumers differ across state years only in a vertical taste-shifter for guns.

4.1 Consumer Choice

We define a market t as a U.S. state $\times$ year. Within each state, the market size M_t is defined to be the fraction of predicted potential gun owners based on the Pew survey on gun ownership described in Section 2, multiplied by the 2015-2019 ACS estimates of adult population in each state. Note that market size is assumed to be constant across years.

We model preferences of consumer i with demographics d in year y. Preferences are over guns j within each weapon class $c(j) \in \{\text{handgun,long gun}\}$. We refer to the set of guns in class c as $\mathcal{J}_{c,y}$, and the entire set of guns on the market as $\mathcal{J}_y = \cup_c \mathcal{J}_{c,y}$. We define $\mathcal{J}_y$ to include all actively produced (those with MSRP data) guns with at least 10 purchases in a particular year y in the FRB, totalling 850 unique models. We use the FRB transaction data from 2016-2022.

Indirect utility from purchasing gun j is defined as follows:

$$u_{i,j,t} = \beta_{i,c} X_j + \delta_{j,t} + \epsilon_{i,j,t}. \tag{5}$$

$$\delta_{j,t} = \alpha p_{j,y} + \delta_j + \tau_t + \phi_{c,y} + \xi_{j,y} \tag{6}$$

$$\beta_{i,c} = \beta_c + D_i \Pi_c \tag{7}$$

where X_j denotes gun j's characteristics, $\delta_{j,t}$ denotes the mean utility of a gun in market t, τ_t denotes market-level shifters in taste, $\phi_{c,y}$ denotes weapon class-by-year demand shocks, and δ_j denotes the time invariant taste for gun j. Note that we normalize $\tau_t = 0$ each year during our demand estimation, so that values are understood relative to Massachusetts, since we only observe transaction-level data in one state. Therefore, mean utilities for estimation can be equivalently expressed as $\delta_{j,y}$. β_c denotes the baseline taste for gun characteristics . D_i is a matrix of consumer demographics, and Π_c is a class-specific matrix mapping demographics to gun characteristic preferences. The price of gun j in year y is denoted $p_{j,y}$, and is the

(CPI-adjusted) MSRP. The price sensitivity of consumers is determined by α.

Though we take a characteristic based modeling approach for the actively produced guns in each market, a substantial fraction of guns purchased in Massachusetts each are used/not actively produced, due to guns being a durable good. These guns may vary from the guns purchased "brand new" that we explicitly model in $\mathcal{J}$, due to variation in condition and price. For this reason, we also include in choice sets a composite good $\omega_{c,y}$, varying at the class and year level, that represents a choice of any used guns purchased that year, guns that are purchased at a very low frequency (< 10), or guns which we are unable to match to a BBGV gun model.[27] The utility the consumer derives from this product is:

$$u_{i,\omega,c,t} = \delta_{\omega,c,y} + \epsilon_{i,\omega,t} \tag{8}$$

Where $\delta_{\omega,c,y}$ includes all utility components found for new guns in Equation 6, except for price. Equivalently, we assume the price of used guns is $p_{\omega,c,y} = 0$ and all price-specific tastes for the composite good are loaded in the unobserved demand shock, $\xi_{\omega,c,y}$.

We assume that the idiosyncratic error term ϵ follows the distribution of a four-level nested logit [Train, 2009], to allow for unobserved correlation between guns. We place the outside option into a separate nest from the inside goods. The inside good nesting parameter is ρ_0. We then partition the inside goods into 2 nests based on weapon class c. The parameters that govern substitutability across the class nests are $\rho_c \leq \rho_0$. Finally, because we do not explicitly model the individual used gun models found in the composite good, we also place the composite good $\omega_{c,y}$ in a separate nest from new products $\mathcal{J}_{c,y}$ within each class, to capture non-independent substitution from used guns to individual guns that are sold as new. The parameter controlling this substitution is denoted $\rho_c^u \leq \rho_c$. ϵ is assumed to have the following structure:

$$\epsilon_{i,j,t} = \zeta_{i,1,t} + \rho_0 \zeta_{i,c,t} + \rho_c \zeta_{i,c,u,t} + \rho_c^u \tilde{\epsilon}_{i,j,t}$$

Values of $\rho \approx 0$ imply a high degree of correlation within a nest, while $\rho \approx 1$ implies no correlation. ζ denote idiosyncratic shocks that are common within each nest, that are conjugate to the extreme value type I distribution [Cardell, 1997], and $\tilde{\epsilon}_{i,j,t}$ is a standard i.i.d extreme-value type 1 shock.

Finally, we specify the value of our outside option as

$$u_{i,0,t} = \tilde{\epsilon}_{i,0,t}. \tag{9}$$

[27] In practice, the vast majority of guns included in the composite good are guns not actively produced in a particular year.

Demographics D_i are composed of a binary gender variable, and characteristics of the zip code the consumer resides in. These zip code characteristics are the % White, % Conservative (as measured by the zip-code % voting for Donald Trump in 2016), the % with a Bachelor's degree or higher, the logged median income of the zip code, and the logged population density per square mile. For each zip code demographic, we center these at the national mean, so that the coefficient β_c can be understood as the taste for a characteristic for a consumer with average demographics. For characteristics X_j, we include caliber, barrel length, a dummy for high-capacity, a constant term (to capture heterogeneous tastes for the outside good), and a dummy for the gun being new, or in $\mathcal{J}_y$ (to capture heterogeneous tastes for used guns).

For notational convenience, it will be useful to write the utility of a consumer i with demographics $d = D_i$ from choosing product j in market t as a non-random component $v_{d,j,t}$ plus the idiosyncratic shock:

$$\begin{aligned} u_{i,j,t} &= v_{d,j,t} + \epsilon_{i,j,t} \\ v_{d,j,t} &= \delta_{j,t} + \beta_{d,c(j)} X_j. \end{aligned}$$

Given the nesting structure specified for $\epsilon_{i,j,t}$, we can express the probability of a consumer i with demographics $d = D_i$ purchasing an individual product j as the product of four nest-level probabilities.

$$Pr(j|d,t) = Pr(j \neq 0|d,t) \cdot Pr(c|d,t,j \neq 0) \cdot Pr(j \in \mathcal{J}_{c,y}|d,t,c) \cdot Pr(j|d,t,j \in \mathcal{J}_{c,y}) \tag{10}$$

$$Pr(\omega_{c,y}|d,t) = Pr(j \neq 0|d,t) \cdot Pr(c|d,t,j \neq 0) \cdot Pr(\omega|d,t,c) \tag{11}$$

$$\tag{12}$$

Within a class c, excluding the composite good, and integrating out the idiosyncratic shock $\tilde{\epsilon}_{i,j,t}$, the probability of consumer i choosing a particular gun model j conditional on buying a new gun $j \in \mathcal{J}_c$ in market t takes the following logit form:

$$Pr(j|d,t,j \in \mathcal{J}_{c,y}) = \frac{\exp\left(\frac{v_{d,j,t}}{\rho_c^u}\right)}{\sum_{k \in \mathcal{J}_c} \exp\left(\frac{v_{d,k,t}}{\rho_c^u}\right)}. \tag{13}$$

As is typical in nested logit models, we proceed by calculating the inclusive value of each nest, then using it to express the choice probability in the next highest nest. The inclusive value, $I_{d,c,n,t}$ for new guns of class c in market t is:

$$I_{d,c,n,t} = \log\left(\sum_{j\in\mathcal{J}_c} \exp\left(\frac{v_{d,j,t}}{\rho_c^u}\right)\right). \tag{14}$$

Moving up the nesting structure, to the choice between used vs new guns within class, the conditional probabilities are:

$$Pr(\omega_{c,y}|d,t,c) = \frac{\exp\left(v_{\omega,c,t}\right)}{\exp\left(v_{\omega,c,t}\right) + \exp(\rho_c^u I_{d,c,n,t})} \tag{15}$$

$$Pr(j \in \mathcal{J}_{c,y}|d,t,c) = \frac{\exp(\rho_c^u I_{d,c,n,t})}{\exp\left(v_{\omega,c,t}\right) + \exp(\rho_c^u I_{d,c,n,t})}1 \tag{16}$$

From which we may derive the class-nest level inclusive values $I_{d,c,t}$:

$$I_{d,c,t} = \log(\exp\left(\frac{\rho_c^u}{\rho_c} v_{\omega,c,t}\right) + \exp(\frac{\rho_c^u}{\rho_c} I_{d,c,n,t})) \tag{17}$$

The corresponding class-level choice probabilities, conditional on not choosing the outside option, are:

$$Pr(c|d,t,j\neq 0) = \frac{\exp(\rho_c I_{d,c,t})}{\sum_{c'\in\{\text{handgun,long gun}\}} \exp(\rho_{c'} I_{d,c',t})} \tag{18}$$

Finally, this yields an inclusive value for all guns, or the inside good:

$$I_{d,1,t} = \log(\sum_c \exp(\frac{\rho_c}{\rho_0} I_{d,c,t})) \tag{19}$$

Which generates the choice probabilities for the inside versus outside option:

$$Pr(j\neq 0|d,t) = \frac{\exp(\rho_0 I_{d,1,t})}{1+\exp(\rho_0 I_{d,1,t})} \tag{20}$$

$$Pr(j = 0|d,t) = \frac{1}{1+\exp(\rho_0 I_{d,1,t})} \tag{21}$$

completing our characterization of the choice probabilities.

The log of the denominator in these top-level choice probabilities plays a role analogous to a standard logit choice model. The quantity $\log(1+\rho_0 I_{d,1,t})$ represents the expected value of the best item in the choice set for a consumer in demographic group d and market t (Train [2009]). We define this value as consumer surplus and convert from utils to dollars by rescaling using the inverse of the price coefficient $1/\alpha$. In this sense, our model of demand allows for flexible cross-good substitution patterns while maintaining a well-defined interpretation of consumer surplus.

4.2 Estimation Routine

We estimate the demand model in five steps. First,using, we estimate the gun model choice multinomial logits within class, to recover $\delta_{j,y}$, Π_c, and ρ_c^u. We then jointly estimate the outside-option and class-choice models, taking the within class nest estimates as given, to recover ρ_c, ρ_0, and $\phi_{c,y}$. We then estimate α, the price parameter, using the instrumental variable for price from Section 3.2. Next, we estimate the baseline parameters governing tastes, β_c, for characteristics via a projection of δ_j on the time-invariant characteristics X_j. Finally, we estimate the market-specific tastes τ_t across state-years using the NICS background check data.

4.2.1 Gun Choice

For each c, we solve the gun choice problem, conditional on choosing class c, to estimate the parameters governing within-class choice, denoted θ_c. We collapse the Massachusetts FRB data to the gender-zip demographic cell level, denoted d, by year y. Let $N_{y,d,j}$ denote the number of individuals purchasing product j in demographic cell d, and $\mathcal{D}_{MA}$ the set of these demographic cells. We solve the following constrained maximum likelihood problem for each class c, as in Train [2009]:

$$\max_{\theta_c} \sum_{y} \sum_{d \in \mathcal{D}_{MA}} \Big(N_{d,\omega} Pr(\omega|d,c) + \sum_{j \in \mathcal{J}_c} N_{d,j} \log(Pr(j|d,c)) \Big) \tag{22}$$

$$\text{Subject to: } Pr(j|y,c) = \hat{s}_{j,y,c} \quad \forall y, j \in \mathcal{J}_c \tag{23}$$

$$\frac{\rho_c^u}{\rho_c} \leq 1 \tag{24}$$

Where the product choice conditional on class is
$Pr(j|d,c) = Pr(j \in \mathcal{J}_{c,y}|d,t(d),c) \cdot Pr(j|d,t(d), j \in \mathcal{J}_{c,y})$ as defined in Equations 13 and 16. The class-level parameters θ_c are the used vs new good nesting parameter ρ_c^u, the product-level taste shifters $\delta_{j,y}$, and the heterogeneous tastes for Π_c, excluding the constant term.[28]

Equation 23 constrains the predicted state-level choice probabilities $Pr(j|y,c)$ to match the empirically observed choice probabilities $\hat{s}_{j,y,c}$ in the FRB data. Given θ_c, $Pr(j|y,c)$ can

[28] Since the constant term is identical the within-class choices $\mathcal{J}_{c,y} \cup \{\omega_{c,y}\}$, we estimate this dimension of preference heterogeneity in the class choice model that includes the outside option.

be calculated by integrating over the distribution of demographic cells d in Massachusetts.

$$Pr(j|y,c) = \int_d Pr(j|d,c))\partial F(d)$$

where $F(d)$ is the empirical distribution of demographic cells. The empirical analogue to this model-derived probability is the quantity $\hat{s}_{j,y,c}$, the empirically observed probability of gun j being chosen in year y, conditional on a firearm from class c purchased:

$$\hat{s}_{j,y,c} = \frac{\sum_{d\in\mathcal{D}_{MA}} N_{y,d,j}}{N_{y,d,\omega,c} + \sum_{j\in\mathcal{J}_{c,y}} \sum_{d\in\mathcal{D}_{MA}} N_{y,d,j}}$$

This constraint exactly identifies the mean utilities $\delta_{j,y}$ (relative to the mean utility for the composite good, which is normalized to zero each year). We implement this constraint via the contraction mapping suggested in Conlon and Gortmaker [2020].

We also impose the constraint that $\rho_c^u/\rho_c \leq 1$ to guarantee the recovered used vs new nesting parameter is consistent with the theoretical model. Because level of ρ_c is not identified in this sub-problem, we estimate the ratio of the two terms directly; this preserves the choice probabilities described in the prior section.[29] Among products within a class, we use the differential buying patterns of consumers of different gender and zip codes to identify the parameters governing the heterogeneous taste for characteristics, Π_c. We use variation in the choice set of $\mathcal{J}_{c,y}$, and it's correlation with the change in the composite good's conditional market share to identify ρ_c^u.

4.2.2 Class Choice

Given estimates θ_c, we calculate the class-level inclusive values, and use these to estimate the parameters θ_0 governing choices across gun classes, and the outside option. These parameters are the class-year demand shocks $\phi_{y,c}$, the heterogeneous preferences for classes, and the nesting parameters ρ_c, ρ_u. We set up a similar constrained maximum likelihood problem to the gun choice model to estimate these parameters:

$$\max_{\theta_0} \sum_y \sum_{d\in\mathcal{D}_{MA}} \Big(N_{d,0,y} Pr(j\neq 0|d) + \sum_{c\in\{\text{handgun,long gun}\}} N_{d,c,y}\log(Pr(c|d))\Big) \quad (25)$$

$$\text{Subject to: } Pr(c|y) = \hat{s}_{y,c} \quad \forall y, c\in\{\text{handgun,long gun}\} \quad (26)$$

$$\rho_c \leq \rho_0 \leq 1 \quad \forall c\in\{\text{handgun,long gun}\} \quad (27)$$

[29] After estimating ρ_c in the class-choice section, we multiply all parameters by ρ_c to ensure both the class and gun choice parameters are on the same scale and consistent with each other.

Where the class-level choice probabilities $Pr(c|d) = Pr(j \neq 0|d) \cdot Pr(c|j \neq 0, d)$ are given by Equations 18 and 20, and $Pr(j \neq 0|d)$ is given by Equation 21. $N_{d,c,y}$ is defined analogously to the gun choice section as the number of individuals observed buying a particular class, and $N_{0,d,y} \equiv M_d - N_{d,y}$, the market size of potential gun consumers in a demographic cell, minus the total number of observed gun purchases in a year. The first set of constraints in Equation 26 enforce that the model-predicted class-level choice probabilities in the state of Massachusetts, $Pr(c|y)$, match the empirically observed probabilities $\hat{s}_{y,c}$. These exactly identify the class-year shocks $\phi_{y,c}$.[30] The second set of constraints ensure we estimate nesting parameters that are consistent with the hierarchical structure of the nests. Heterogeneous tastes for classes are identified by the differential propensities of individuals of with differing gender and zip code characteristics to purchase handguns, long guns, or the outside option. The correlation between $\mathcal{J}_{c,y}$, which shifts the class-level inclusive values $I_{c,y}$, identifies ρ_c. Finally, the correlation between changes to the size of $\mathcal{J}_y$ and the outside good identify ρ_0.

4.2.3 Additional Estimation Steps

Having recovered the mean utilities $\delta_{j,t}$ from the gun and class choice models, we use these quantities to identify the utility components that do not vary across individuals. First, we estimate δ_j and α using the following 2SLS regression with two-way fixed effects:

$$p_{j,t} = \delta_{j,1} + \psi_{c,y,1} + \sum_{k=1}^{3}(\beta_{m(j)}\tilde{p}_{k,t}) + \epsilon_{j,t} \tag{28}$$

$$\delta_{j,y} = \delta_j + \psi_{c,y} + \alpha\hat{p}_{j,t} + \xi_{j,t} \tag{29}$$

Where we use the same instrumental variables as in Section 3.2, and the same LASSOIV procedure, to identify the price coefficient α. Note this regression excludes the composite good, which has no observed price, as well as guns with zero sales, as they imply a value of $\delta_{j,t} = -\infty$.

Next we take our estimates of unobserved, time-invariant differences in gun models, $\hat{\delta}_j$, and project these onto the time-invariant characteristics X_j to construct estimates of β_c. We do so using the approach outlined in Nevo [2000], which weights the mean tastes δ_j by the estimated variance of the fixed effects.

Finally, we use the NICS background check data to identify τ_t. Explicitly, we assume that the number of state-year background checks are equal to the inside good share $\hat{s}_{1,t}$ in each market t. Since τ_t affects all goods equally in any market, it can be separated from the

[30]These class year shocks $\phi_{y,c}$ are then added to the gun choice estimates of product-level mean utilities, $\hat{\delta}_{j,t}$, to get our final estimates of the mean utilities, e.g. $\delta_{j,t} = \phi_{y(t),c(j)} + \rho_{c(j)} \cdot \hat{\delta}_{j,t}$

inclusive value of the inside good. We solve for the value of τ_t in each market that sets the model-predicted choice probabilities to their observed inside good share:

$$\tau_t : \quad Pr(j \neq 0|t) = \int_d \frac{\exp(\tau_t + \rho_0 I_{d,1,t})}{1 + \exp(\tau_t + \rho_0 I_{d,1,t})} F_t(d) = \hat{s}_{1,t}$$

where $F_t(d)$ is the distribution of zip code characteristics $\times$ gender in market t.

4.3 Assumption and Limitations

The above demand model is able to flexibly estimate heterogeneous demand by consumer types, various sources of unobservable utility, as well as account for non-independent substitution patterns across weapon products. Moreover, casting the problem as a discrete choice problem allows for tractable estimation. However, there are some key assumptions we make in the model we wish to highlight here. First, we model the gun purchasing decision as a discrete choice problem, treating observed purchases as effectively the "share" of consumers in a market who purchase a firearm. However, it is a well known fact that most gun purchases are from those who are repeat buyers and already own a weapon [Miller et al., 2022]. Because we lack data on individual identifiers in the FRB data, we cannot disentangle these two types of consumers, who may differ in their preferences. Therefore, we interpret our estimated demand parameters as a population-weighted blend of extensive and intensive margin preferences for firearms.

Second, we constrain α, the price sensitivity of the consumer, to be homogeneous across the population. This is due to our lack of demographic-level variation in exogeneous price shocks. In their conjoint analysis, Moshary et al. [forthcoming] allow for the price coefficient to vary along demographics, and show, at least for income and age, statistically significant differences in price heterogeneity, although they claim that this price heterogeneity is not particularly important for their analysis.

Our third and strongest assumption is that the differences in demand for firearms differ only demographics and the vertical market taste shifter τ_t. We only have access to transaction-level data for the state of Massachusetts, so our parameters (besides τ) are entirely identified by demand variation within that state. Massachusetts is not a typical state with respect to its attitudes towards firearms: it one of the most restrictive states in terms of gun control[31]. For a handguns to be legally sold to consumers, they must be registered on the state's approved firearms roster,[32] which impacts the choice sets $\mathcal{J}_{c,y}$ we infer are

[31]Source: https://www.rand.org/research/gun-policy/firearm-mortality.html

[32]See https://www.mass.gov/lists/approved-firearms-rosters. Long guns are under no such restrictions in Massachusetts.

available to consumer's in other state in a particular year. Because of this, we focus much of our counterfactual exercises on policy changes occurring solely in the state of Massachusetts. At the same time, we also provide demand-side estimates of the impacts of these policies in the entire United States, for those readers willing to extrapolate our estimated demand parameters to the rest of the country, but these estimates should be viewed more cautiously.

4.4 Demand Estimates

We now describe our demand estimates. In Table 4, we report the nesting parameters, along with p-values for the hypothesis test that there are correlated intra-nest preferences (e.g. for ρ_c, the null hypothesis is that $\rho_c = \rho_0$, and the alternative $\rho_c < \rho_0$). The outside option displays a nesting parameter of $\rho_0 = .9$, suggesting reasonably flexiblity between buying a gun and choosing to purchase nothing. However, this is eclipsed by the class-level nesting parameters, ρ_c which are estimated to be around 0.3. These imply that consumers have strong preferences to substitute only to other products that are the same class of weapon; this tendency is slightly less so for long gun consumers, meaning they are more likely to view handguns as acceptable substitutes to their preferred firearm class. Regarding the used vs new good substitution, we estimate that there is reasonably flexible substitution between the used composite good and new firearms: we cannot reject the hypothesis that $\rho_c^u = \rho_c$ at conventional statistical levels. This is consistent with anecdotal evidence that during COVID 19, many gun manufacturers reached binding capacity constraints and were unable to keep up with demand,[33] leading to the surge in used gun purchases we document in Figure 5.

In Appendix Table A1, we report the estimate of the price coefficient from estimating the 2SLS system in Equations 28 and 29.[34] The coefficient is stable across specifications. Our preferred estimate is located in Column (2), with class $\times$ year fixed effects. We use this coefficient to estimate the price responsiveness of firearms consumers, as well as to monetize the remaining model parameters.

Panel (a) of Figure 10 displays the estimated own price elasticities at the gun-product level in the state of Massachusetts.[35] We estimate an average (median) own-price elasticity of -2.2 (-1.7), consistent with our reduced form results. Elasticities differ somewhat across firearm class: handgun consumers are slightly more price elastic than long gun consumers (-2.3 vs -2.0), but for both types of firearms, there is a long left tail of models with highly

[33] See for example `https://cbsaustin.com/unprecedented-demand-on-guns-and-ammo-putting-pressure-on-supply-chain`

[34] Because this regression is in price levels, we remove one single gun model from Barrett Firearms (5 observations), which has large swings in price and reaches over $12,000, and has excess influence over our price parameter estimates. Removing other gun models has no qualitative effect on the parameter estimate.

[35] In both panels, estimates are very similar for the entire United States.

elastic demand.

Panel (b) shows how consumers substitute across products classes. In particular, it presents the average diversion ratios [Conlon and Mortimer, 2021] from handgun and long guns to weapon classes and the outside option. Diversion ratios are a statistic that represent the following thought experiment: given a price increase in product j, what will the consumers who choose to no longer purchase product j due to the price increase choose instead? Our low estimates of the class-level nesting parameter ρ_c imply that the consumers have strong intra-class preferences: for example, the vast majority of handgun consumers will switch to another handgun when the price of their first-choice product increases. At the same time, there is a sizeable mass of consumers who substitute to the outside option and choose to purchase no firearm at all.

In Table 5, we present estimates of β_c and Π_c in money-metric terms, scaling our estimates by the price coefficient α. This allows us to interpret parameter estimates in terms of a consumer's willingness to pay (in dollars) for a product that differs in one characteristics but is otherwise identical to the average product in the market. The full table of estimated utility parameters, with standard errors, can be found in Appendix Table A2. The large estimated coefficients on the constant term suggest most consumer heterogeneity loads onto whether the firearm is a handgun or long gun, which serve distinct market segments. For example, in the Pew ATP survey, 72% of handgun owners say protection is a major reason they own a gun, while for long gun owners, it is only 35% (and instead, hunting is the most popular reason). Among the most significant drivers of weapon class preferences are gender, political leaning, racial composition, and education levels of neighborhoods (zip codes). Regarding individual gun characteristics, we find that females display strong preferences for smaller guns. Those living in Republican areas prefer smaller weapons with higher capacity and caliber, while those in whiter areas prefer larger weapons. Consumers from more affluent neighborhoods have relatively higher tastes for shotguns and high capacity firearms.

In Appendix Figure A5, we display the average estimate of the cross-market shifter τ_t/α for each state during our sample period. As the model perfectly predicts firearms sales in Massachusetts $\tau_{MA,y} \equiv 0$, by construction, this quantity captures additional preference for firearm purchase relative to Massachusetts, after accounting for differing demographics across the U.S. The figure shows considerable variation in taste for firearms, ranging in the thousands of dollars. Considering Massachusetts has one of the lowest inside good shares in the country (3.7%), it is unsurprising that it's value is low relative to other states. Some interesting patterns emerge from the figure. For example, California has an inside share of 4.2%, similar to Massachusetts. Yet California's consumers are willing to pay $1,200 more for the same firearms, conditional on demographics. This stems from the preference estimates in Table 5: California is a racially diverse and politically left-leaning state, so the demographic preferences predict that demand should be quite low. The fact that we see even more individuals buy firearms in California than in Massachusetts implies that Californians must have unobservably higher preferences for firearms.

5 A Model of Public Health

Although firearms connect to a wide array of outcomes across society [Everytown, 2022], many of these are determined through potential incidents of firearm violence [Lott and Mustard, 1997, Duggan, 2001]. Recent critical reviews have noted the potential for bias in existing estimates of the effect of firearms on violence, due to errors in the measurement of firearms Kleck [2015] and the potential for reverse causality [RAND, 2018]. In this section, we provide an estimate of the causal effect of licit firearm purchases on downstream incidents of firearm violence. To do so, we measure fatal and non-fatal firearm violence using the high-resolution data from the Gun Violence Archive. We instrument for firearm purchases in high resolution using the entry of local firearm retailers, as proposed and implemented by Rosenberg [2023]. We will use these estimates to quantify the net welfare effects of firearm policy counterfactuals, accounting for the downstream effects of purchases on public health.

5.1 Model Specification

We propose to model gun violence incidence in five-digit zip code tabulation area (ZCTA5) z during year t as a Poisson random variable, whose mean depends on contemporaneous legal firearm transactions and two-way fixed effects:

$$E[\#\text{GunViolence}_{z,t}|\#\text{GunTransactions}_{z,t}] = \exp(\theta \log(1 + \#\text{GunTransactions}_{z,t}) + \psi_z + \phi_{s,t}) \tag{30}$$

$$\#\text{GunViolence}_{z,t} = E[\#\text{GunViolence}_{z,t}|\#\text{GunTransactions}_{z,t}] \cdot \omega_{z,t}. \tag{31}$$

Where $\#\text{GunTransactions}_{zt}$ denotes the number of gun transactions in zip code z during year t, as recorded by the Masscusetts FRB. We employ a Poisson model to handle areas with zero deaths, and for the handful of areas with no firearm transactions, we utilize the $\log(1 + \cdot)$ transformation for gun transactions. This regression specification further includes two-way fixed effects for zip code ψ_z and state-year $\phi_{s,t}$, the latter allowing us to control for changes in state-level firearms policies. We treat the multiplicative regression errors ω_{zt} in Equation (31) as idiosyncratic violence shocks that capture other causes of gun violence aside from contemporaneous gun purchases (e.g., pre-existing firearm owners, transactions on the illicit market, transactions in other geographies).

Evaluating the model at zero firearm transactions defines a baseline level of gun violence

$$\#\text{GunViolence}_{z,t}\Big|_{\#\text{GunTransactions}_{z,t}=0} = \exp\left(\psi_z + \phi_{s,t}\right) \cdot \omega_{z,t}$$

From this baseline, the model implies constant elasticity effect of θ from additional firearm transactions.

5.2 Estimation and Instruments

Two factors complicate estimation of the regression in Equation (31). The first issue is the potential for endogeneity. In particular, the unobservable shocks to gun violence ω_{zt} may be correlated with shocks for firearm demand. This is a natural concern, given that studies have shown individuals change their firearm registration behavior in response to episodes of local violence [Depew and Swensen, 2019], and personal protection is the most commonly stated reason for gun ownership [Parker et al., 2017].

To overcome the issue of endogeneity, we leverage the entry and exit of firearm retailers into local markets as shifters of gun demand. These instruments are developed by Rosenberg [2023], who discusses the local retail market for firearms at length and conducts a number of tests for instrument relevance and validity.

We specify our instrument for firearm transactions as the natural log of one plus the number of firearm retailers in a local market.

$$w_{z,t} = \log\left(1 + \text{FirearmRetailers}_{z,t}\right).$$

Which yields the following first-stage regression equation:

$$\log\left(1 + \#\text{GunTransactions}_{z,t}\right) = \beta w_{z,t} + \psi_z^{fs} + \phi_{s,t}^{fs} + \epsilon_{z,t} \tag{32}$$

We use a national measure of local firearm retailer presence from the Historical Business dataset provided by Data Axle (formerly InfoGroup)[36] This proprietary dataset provides a snapshot at the end of each calendar year of local businesses in operation, compiled via yellow page data across the United States. Each business is associated with geographic identifiers, such as ZCTA5, latitude/longitude, and an extended 8-digit NAICS code for

[36] Data was acccessed via Wharton Data Research Services on December 17th, 2023. See more information about the dataset at `https://wrds-www.wharton.upenn.edu/pages/support/manuals-and-overviews/infogroup/wrds-overview-infogroup/`

business type classification. We download all business location records from 2014 to 2022 for those businesses with an extended NAICS code of 451110-23, identified by Data Axle as retail gun shops (under the standard 6 digit NAICS code of 451110, sporting goods stores). We then aggregate these gun retailer records to the ZCTA5-year level to construct our instrument w_{zt}.

The key exclusion restriction is that firearm retailers do not vary the timing of their entry and exit decisions in response to local variation in gun violence, so that net entry impacts violence only through firearm sales. In addition to extensive empirical tests, Rosenberg [2023], discusses the *a priori* reasonable nature of this assumption. The highly regulated nature of the retail firearm market makes it challenging for potential entrants to time their entry decisions based on the anticipation of local violence shocks. We validate this claim in the next subsection by showing that, conditional on the contemporaneous market structure, firearm sales are uncorrelated with lagged entries and exits of firearm retailers within a zip code.

The second issue in estimation is that our FRB data on transactions, the endogeneous variable of interest, only covers one small state within the U.S. To fully utilize our national data on gun violence incidents and firearm retailers, we estimate Equation (31) using nonlinear two sample IV, similar to the approach in Hastings et al. [2021]. Explicitly, we estimate Equation 32 using the FRB data in Massachusetts (2016-2022), to recover the first-stage coefficient β. Then, we predict the number of transactions in each zip code-year across the United States attributable to the net entry of firearm retailers, $\log(1+\#\text{Gun}\widehat{\text{Transactions}}_{zt}) = \beta \cdot w_{z,t}$. We substitute this prediction for the endogeneous (and unobserved) actual transactions in each zip-year from 2016 to 2022 when estimating the causal effect of firearm transactions on gun violence, θ, via a Poisson regression of Equation (31).[37] For inference, we bootstrap this procedure 1,000 times, resampling at the zip code level, and stratifying our sampling at the state level (to ensure Massachusetts is appropriately covered in each replicate).

5.3 Results

Figure 11 visualizes the first-stage effect of our net entry instrument against two alternative measures of gun retail presence in Massachusetts. The first uses the FRB data to reconstruct the set of firearm retailers that were open within each zip code-year.[38] Although this method is accurate, covering the universe of legal firearm retailers, it can only be constructed for

[37] For this approach to be valid, we assume that the estimate of β in Massachusetts is representative for the rest of the United States.

[38] Specifically, we use the last and first year retailers had at least one transaction to determine whether they are in the market. If a retailer opens in the middle of the year, we assign the instrument a value proportional to the months of the year they were present.

Massachusetts, and cannot be used for national-level estimates. The second measure follows the literature and measures firearm retailers as the count of Federal Firearms Licensees ("FFLs"), representing all entities who are licensed to sell firearms [Wiebe et al., 2009, Johnson and Robinson, 2021]. This measure can be constructed at the national level but is much noisier, since many FFLs are "kitchen-table" dealers and personal collectors who are in fact unavailable to many firearms consumers.[39]

Figure 11 displays a best-fit line from the first stage equation along with the optimal number of bins for the relationship as determined by the data-driven method of Cattaneo et al. [2019] across these three measures of the firearm retail market. Both the FRB and Infogroup measures positively covary with firearms transactions, and we estimate a very similar elasticities, suggesting that Data Axle is a reasonable proxy for the universe of firearm retailers in Massachusetts. On the other hand, the relationship between sales and FFLs is noisy and statistically indistinguishable from zero, highlighting challenges with this measure of the local firearm market.

To ensure that our first-stage estimates are not driven by differential pre-trends (demand for guns increasing and retailers responding to anticipated demand via entry), we also estimate a dynamic version of Equation (32). In this specification, we add predictor variables using retailers in years $t-1$ and $t-2$ in each zip code. These estimates are displayed in Appendix Figure A4. Conditional on the contemporaneous firearm market, we find no evidence of positive covariation between past entries of firearm retailers and contemperaneous firearms sales. This suggests that the net entry of firearm retailers is uncorrelated with unobservable determinants of local firearm demand, supporting our key exclusion restriction.

Table 6 presents estimates of the effect of firearms sales on gun violence. We use three measures of gun violence: number of non-fatal injuries, number of fatal injuries, and number of incidents (regardless of the number of fatal or non-fatal injuries per incident). In Panel A, we display estimates of the elasticity of firearm violence with respect to retailer net entry (the intent-to-treat effect), which can be estimated via a Poisson regression. The estimates are all fairly precise and suggest that retailers have a greater effect on injuries than deaths.

In Panel B of Table 6, we display our preferred two-sample IV estimates for the effect of firearm transactions on gun violence. Across the board, higher volumes of firearm transactions induced by the net entry of firearm retailers are estimated to cause contemporaneous increases in firearm-related injuries, and fatalities. We estimate that a 10% increase in firearm transactions would increase gun fatalities by 4%, non-fatal gun injuries by 10%, and incidents of firearm violence by 9%. For inference, we report bias-corrected bootstrap

[39] A regression of Massachusetts retailers on FFLs has a coefficient of approximately 0.3,imply that only 1-in-3 FFLs are actively sell to consumers.

confidence intervals [Efron, 1987], since our bootstrap estimates exhibit right skewness. The estimated magnitudes in these transaction-based regressions are comparable in size to the estimated causal effects of gun ownership on gun violence in prior studies [Duggan, 2001, Cook and Ludwig, 2006, Rosenberg, 2023].

5.4 Using the model

The estimates from our constant elasticity model are easy to use for counterfactuals. For example, suppose we contemplate a counterfactual that *ceterus paribus* shifts gun transactions in one zip code-year from g_{zt} to g'_{zt}. Using Equation 31, the expected change in gun violence is

$$\#\text{GunViolence}_{zt}\Big(g'_{zt}\Big) - \#\text{GunViolence}_{zt}\Big(g_{zt}\Big) = \#\text{GunViolence}_{zt}\Big(g_{zt}\Big) \times \left\{ \left(\frac{1+g'_{zt}}{1+g_{zt}}\right)^{\theta} - 1 \right\}. \tag{33}$$

This quantity depends on observed levels of violence and both observed and and counterfactual firearm transactions, while the unobservable components of the model difference out. By evaluating the right-hand side, we can recover the implied effects on gun violence of any specified change in gun transactions. Summing these differences across time and space provides a natural method of aggregating public health effects under the model.

Less clear is how to convert our results across different outcome variables into money-metric welfare measures. To do so, we utilize calibrations from the criminology literature that value each life lost due to homicide at 9.5 million dollars and each non-fatal assault at 92 thousand dollars [Peterson et al., 2021].[40] Although these are challenging quantities to estimate, we view them as a useful tool for summarizing the disparate consequences of firearm regulation on a single, interpretable scale. We also note that these estimates are relatively conservative, since gun violence incidents tend to skew towards younger individuals, which implies a greater loss in the quality of life.[41]

6 Counterfactuals

In this section, we synthesize our empirical results to study the impacts of counterfactual regulations on the market for consumer firearms. We use our estimates of preferences for

[40] Calculated as the total cost from loss of quality of life, medical expenses, and lost productivity, divided by number of incidents, and adjusted to 2022 dollars via the CPI.

[41] For example [Everytown, 2022] estimates a cost of about 15 million dollars per gun homicide.

gun purchase to predict the effects of policies on consumer surplus, government revenues, and the public health. We carry out this exercise twice, once focusing only on Massachusetts as the source of our transaction data and a second time extrapolating our demand estimates from Massachusetts to the entire United States.

In our counterfactuals, we do not allow firearm manufacturers to adjust their decisions on the supply-side of the market in response to policy. This is a good approximation of the likely supply-side effects from regulation specific to Massachusetts, which makes up only a small fraction of the national firearms market. However, it likely provides a poor guide to the effects of national-level policy. In addition, it requires extrapolation of the demand model estimated with only Massachusetts data to reflect tastes for firearms across the United States (up to a vertical preference shifter τ across markets). Therefore, we interpret our counterfactual exercises as reasonable approximations for the effects of policy changes in Massachusetts, and a demand-side only analysis of the effect of these policies throughout the United States.

6.1 Mechanics and Implementation

Before presenting counterfactual predictions, we first describe the machinery and auxiliary assumptions we use to compute welfare estimates across our proposed policies.

We first consider policy alternatives based on counterfactual specific additive taxes (e.g. \$100 increase in price) on handguns and long guns. We focuses on additive taxes, as opposed tax rates, because this price increase is well defined for both new gun in $\mathcal{J}_{c,y}$ and the used gun composite $\omega_{c,y}$, which in our demand model does not have a set price. The taxes we consider range from a moderate subsidy- in which we decrease the price of each gun buy 100 dollars- to a 3,000 dollar price increase, representing roughly a 400% increase in the price of guns.

We now describe implementation of the additive taxes, and our framework for welfare evaluation. Let T_j denote a specific tax amount for product j. Since taxes affect consumer utility homogeneously, we implement these taxes in our model by shifting the mean utilities for each product $\delta_{t,j}$:

$$\delta_{j,t}(T_j) = \delta_{j,t} + \alpha T_j$$

and recalculating choice probabilities $s_{j,t}$ from Equation 10 according to the new $\delta_{j,t}(T_j)$. Consumer surplus is calculated in the typical form for discrete choice models

$$CS(\vec{T}) = \sum_t CS_t(T_j) = \sum_t M_t \int_d -\frac{1}{\alpha} \log\left(1 + \exp(\rho_0 I_{d,1,t}(\vec{T}))\right) \partial F_t(d) \tag{34}$$

where t indexes markets, $\vec{T}$ is the vector of taxes across all products, M_t is market size, and

$I_{d,1,t}(\vec{T})$ is the inclusive value of the inside option from Equation 19, recalculated to account for the new mean utilities $\delta_{j,t}(T_j)$.

While consumer surplus will decrease from the policies, there are two potentially offsetting benefits: increases in tax revenue, and increases in public health. For taxes these are simply calculated as:

$$\mathcal{T}(\vec{T}) = MVPF \cdot \sum_t M_t \sum_j T_j \cdot s_{j,t}(\vec{T}) \tag{35}$$

Where $MVPF$ is the marginal value of public funds Hendren and Sprung-Keyser [2020]. We follow O'Connell and Smith [2021] and value each dollar of tax revenue raised by the government as producing one dollar of social welfare, i.e. $MVPF = 1$. Public health benefits are calculated as in Equation 33

$$\mathcal{PH}(\vec{T}) = \sum_i \Psi^i \sum_t \sum_{z \in t} \bar{h}_{i,z,t} \Big(\Big(\frac{1 + \sum_j q_{j,z,t}(\vec{T})}{1 + \sum_j q_{j,z,t}(\vec{0})} \Big)^{\theta^i} - 1 \Big) \tag{36}$$

where i denotes incident types (injury, fatality), Ψ^i denotes the conversions of incident type i to dollars, $\bar{h}_{i,z,t}$ denotes the baseline incidents observed with no taxes, $q_{z,j}$ is predicted purchases of each product in a zip code from the demand model, and θ^i is the incident elasticity estimated in Section 5. As mentioned in Section 5, we follow prior work and value each gun assault fatality at -9.5 million dollars and each non-fatal firearm injury at -0.092 million dollars [Peterson et al., 2021]. As our health model only estimates the effect of gun transactions on contemporaneous firearm violence, we take this short-term view as our baseline measure of externalities.

We calculate aggregate welfare changes from gun taxes as follows:

$$\mathcal{W}(\vec{T}) = \Delta CS(\vec{T}) + \mathcal{T}(\vec{T}) + \mathcal{PH}(\vec{T}) \tag{37}$$

Where ΔCS is the change in consumer surplus.

In a second set of counterfactuals, we consider bans on certain classes of firearms. These are equivalent to taxes $T_j = \infty$ on certain products in the market, or equivalently, setting the inclusive value corresponding to these classes as $I_{d,c,t} = -\infty$, while also setting $\mathcal{T} = 0$.[42]

We consider (jointly) two types of taxes in this paper: taxes on all handguns, and taxes on all long guns. We evaluate counterfactuals under the complete two-dimensional grid spanned by our counterfactual tax rates.

For estimating counterfactuals in the entire United States, we require an additional first-

[42] We implement these numerically by instituting a one million dollar tax on all covered products of the ban; We verify that empirically no consumer purchases any affected products.

step to align the resolution of the data with the resolution of our model. In particular, certain model components—including the health model—are non-linear aggregates of zip code-year level transaction quantities. However, outside Massachusetts, we only observe quantities of firearm transactions at the market (state-year) level. To bridge this gap, we use our fitted demand model to impute firearm transaction quantities under the status quo for each zip code-gender-year outside of Massachusetts, which is feasible since we do have zip-code level data on demographics for the entire United States. The mechanics of this imputation are identical to our predictions within Massachusetts under counterfactual prices.

6.2 Counterfactual Gun Policies

We begin our analysis in Panel A of Figure 12 by examining the value of the firearm market to consumers. In Massachusetts, where we estimate our demand system, we find that the licit firearm market produces approximately 214 million dollars of consumer surplus each year, about 38 dollars per adult per year. This value varies meaningfully over time, reflecting peaks and troughs in demand for gun purchase within the population.

Figure 13 considers how intermediate outcomes, as well as total welfare, change as a function of specific taxes on the purchase of handguns and long guns, on average across years of our sample. In the Figure, we vary the tax from \$0 to \$1,500, which would roughly triple the prices of firearms. In Panel (a), we see that demand slopes downward in our estimated preference model, as sales of each weapon type are decreasing in their respective tax. We estimate that there is non-negligible cross-substitution, in terms of sales, to the type of weapon that is not taxed for the class-specific taxes. However, as our estimated class-level nesting parameters imply, the amount of substitution is quantitatively small. For example, at a \$200 increase in taxes for handguns, approximately 8,600 less handguns are sold per year, but 600 more long guns are sold.

In Panel (b), we plot the various welfare components from Equation 37 as a function of the tax amount. To be of the same sign as other components, consumer surplus is shown in terms of how much it decreases as we raise taxes. The Figure shows that there is a direct tradeoff between consumer surplus and public health, as higher gun taxes worsen market outcomes but prevent gun violence. At the same time, the health benefits of reduced gun sales are an order of magnitude larger than the costs to consumer surplus. Tax revenue follows a Laffer curve, but the price inelastic demand leads to a peak at specific tax values of greater than 1,000 dollars per handgun. Overall welfare benefits are increasing, but concave, in higher taxes.

In Figure 14, we plot the frontiers of positive outcomes (government revenue and public health) against consumer surplus from each type of firearm tax. Implementing the tax rate that maximizes this Laffer curve would require more than tripling the price of the average new handgun in our sample. For public health, We find that the drop in gun violence from tax increases pays for itself in terms of social welfare. The relative length of each frontier in Figure 14 demonstrates a complementarity between taxes on hand guns and long guns. Taxes on individual classes fail to achieve the same magnitude of benefits of a uniform tax due to substitution to other firearms.

Figure 15 expands the policy space to consider jointly varying taxes on both handguns and long guns, by plotting the average welfare benefit per year of a particular tax scheme. We find that raising taxes on handguns or long guns would increase welfare. The two tax instruments work complementarily. The best outcomes from a social welfare perspective are achieved when both taxes are simultaneously high. Setting high taxes on a single weapon class without addressing the other leads to the least socially desirable outcomes. These patterns follow from the challenge of designing policies strong enough to deter gun purchases and decrease gun violence, rather than to merely decrease consumer surplus.

In Figure 16, considers bans on weapon classes, which are the strongest policies available to regulators. We benchmark these bans against a \$100 gun tax, which would represent a $\approx 13\%$ increase in guns, comparable to recently proposed state level gun tax increases,[43]and a \$800 tax, which would roughly double prices. Similar to our tax results, we find that unilateral bans on one weapon class are less effective, due to substitution effects. However, banning both weapon classes- closing the Massachusetts gun market- would drastically reduce gun violence and improve social welfare by 892 million dollars each year. Most importantly, the figure shows that the effect of banning all guns dwarfs the effects of even doubling gun prices (\$267 million increase in welfare), despite the potential government revenue benefits of taxes.

We also conduct our partial equilibrium counterfactuals for the U.S. as a whole. These are shown in Panel (b) of Figures 12,15, and 16.

In Panel B of Figure 12, we consider the spatial distribution of consumer surplus from the firearm market across the U.S. The average adult in the continental U.S. enjoyed 121 dollars of consumer surplus in the average year from 2016-2022. Using a U.S. adult population of approximately 250 million, we value the U.S. firearm market as producing approximately 31 billion dollars of value for consumers each year. However, these gains are unevenly distributed across space. Consumers in urban centers and along the oceanic coasts value the consumer firearms market at roughly one-half of the U.S. population mean. Consumers on the Gulf

[43]For example, California Assembly Bill 28 `https://legiscan.com/CA/text/AB28/id/2842856` and New York Senate Bill 7733 `https://www.nysenate.gov/legislation/bills/2023/S7733`

Coast and in the northern plains value the firearms market at 50 and 100 percent higher than the average, respectively.

In Panel B of Figure 16, we display the effects of shutting down the entire market for consumer firearms in the U.S. This counterfactual is far from the observed data in the U.S., so we interpret the results cautiously. We find that a complete ban on all firearm transactions would increase welfare in the U.S. by 134 billion dollars per year.[44] The reduction in gun violence is far larger than what is achieved through more intermediate tax policy, leading complete bans to dominate taxes from a social welfare perspective.

We find that the welfare gains from gun bans are concentrated in states with racial diversity. Appendix Figure A6 plots the gains (per capita) from a complete firearm sales ban against the share of the population that is white. The gains are approximately flat at \$600 per capita until 60% of the population is white, at which point, welfare gains decrease linearly, to the point states with few racial minorities in the Northeast experience virtually no increase in welfare. This is due to the fact that states with low levels of racial minorities are among the most intensive in terms of firearm purchasing frequency. As a result, eliminating legal firearms transactions has the largest impact on their consumer surplus. This suggests that flexibility and locally tailoring of firearm regulations may be a valuable policy tool. In particular, since demand for guns (and associated public health consequences) vary across populations, it may be that policymakers would be better served targeting the most at-risk jurisdictions with the most stringent policies, rather than a uniform policy across the entire country.

7 Conclusion

In this paper, we provide the first characterization of demand for licit firearms in the United States based on observed transactions. Through descriptive analysis and a structural model, we show that firearms consumers are relatively price inelastic, and differ demographically in their preferences over the characteristic space of firearms in meaningful ways. We then estimate a public health model to quantify the externalities in terms of gun violence that result from increased gun purchases. In our counterfactual policy exercises, we show that even though substitution between types of firearms is relatively low, uniform taxes are more welfare enhancing than targeted taxes. However, due to the limited price responsiveness of these consumers, the largest gains are obtained when we shut down the firearms market

[44] For comparison, Everytown [2022] estimates that gun violence costs the U.S. 557 billion dollars per year; our health estimates from shutting down the market for legal firearms leads to implied health increases of \$165 billion.

entirely. While our paper is an important step towards understanding this market, many questions are left unanswered. In particular, our demand-side analysis fails to account for the supply-side response to taxes, and how this may impact estimates of overall welfare, both through the price response of gun manufacturers, as well as the reduced profits this industry may enjoy. Additionally, this paper does not speak to potential regulatory on the illicit market for firearms, and how this may respond to changes in government policy. Further work is necessarily in the market to better understand the implications of the large scale policies proposed in this paper.

References

Matthew Backus, Christopher Conlon, and Michael Sinkinson. Common ownership and competition in the ready-to-eat cereal industry. Technical report, National Bureau of Economic Research, 2021.

Alexandre Belloni, Daniel Chen, Victor Chernozhukov, and Christian Hansen. Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica*, 80(6):2369–2429, 2012.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

Steven T Berry and Philip A Haile. Foundations of demand estimation. In *Handbook of industrial organization*, volume 4, pages 1–62. Elsevier, 2021.

Sofia Berto Villas-Boas. Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies*, 74(2):625–652, 2007.

Douglas C Bice and David D Hemley. The market for new handguns: an empirical investigation. *The Journal of Law and Economics*, 45(1):251–265, 2002.

Jurgen Brauer. Demand and supply of commercial firearms in the united states. *The Economics of Peace and Security Journal*, 8(1), 2013a.

Jurgen Brauer. The us firearms industry, 2013b.

Timothy F Bresnahan and Peter C Reiss. Entry and competition in concentrated markets. *Journal of political economy*, 99(5):977–1009, 1991.

N Scott Cardell. Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, 13(2):185–213, 1997.

Matias D Cattaneo, Richard K Crump, Max H Farrell, and Yingjie Feng. On binscatter. *arXiv preprint arXiv:1902.09608*, 2019.

Christopher Conlon and Jeff Gortmaker. Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4):1108–1161, 2020.

Christopher Conlon and Julie Holland Mortimer. Empirical properties of diversion ratios. *The RAND Journal of Economics*, 52(4):693–726, 2021.

Philip J Cook and Jens Ludwig. The social costs of gun ownership. *Journal of Public Economics*, 90(1-2):379–391, 2006.

Briggs Depew and Isaac D Swensen. The decision to carry: The effect of crime on concealed-carry applications. *Journal of Human Resources*, 54(4):1121–1153, 2019.

Mark Duggan. More guns, more crime. *Journal of political Economy*, 109(5):1086–1114, 2001.

Bradley Efron. Better bootstrap confidence intervals. *Journal of the American statistical Association*, 82(397):171–185, 1987.

Everytown. The economic cost of gun violence. Technical report, Everytown Research and Policy, 2022. URL `https://everytownresearch.org/report/the-economic-cost-of-gun-violence/`.

Matthew Gentzkow, Jesse M Shapiro, and Matt Taddy. Measuring group differences in high-dimensional choices: method and application to congressional speech. *Econometrica*, 87 (4):1307–1340, 2019.

Paul LE Grieco, Charles Murry, and Ali Yurukoglu. The evolution of market power in the us automobile industry. *The Quarterly Journal of Economics*, page qjad047, 2023.

Justine Hastings, Ryan Kessler, and Jesse M Shapiro. The effect of snap on the composition of purchased foods: Evidence and implications. *American Economic Journal: Economic Policy*, 13(3):277–315, 2021.

Nathaniel Hendren and Ben Sprung-Keyser. A unified welfare analysis of government policies. *The Quarterly Journal of Economics*, 135(3):1209–1318, 2020.

David Blake Johnson and Joshua J Robinson. Gun dealer density and its effect on homicide. *Available at SSRN 3867782*, 2021.

Mark R Joslyn, Donald P Haider-Markel, Michael Baggs, and Andrew Bilbo. Emerging political identities? gun ownership and voting in presidential elections. *Social Science Quarterly*, 98(2):382–396, 2017.

Gary Kleck. The impact of gun ownership rates on crime rates: A methodological review of the evidence. *Journal of criminal justice*, 43(1):40–48, 2015.

Brian Knight. State gun policy and cross-state externalities: Evidence from crime gun tracing. *American Economic Journal: Economic Policy*, 5(4):200–229, 2013.

Christoph Koenig, David Schindler, et al. Dynamics in gun ownership and crime-evidence from the aftermath of sandy hook. Technical report, Working paper, 2016.

Christopher S Koper and Jeffrey A Roth. The impact of the 1994 federal assault weapon ban on gun violence outcomes: an assessment of multiple outcome measures and some lessons for policy evaluation. *Journal of Quantitative Criminology*, 17:33–74, 2001.

Matthew Lang. State firearm sales and criminal activity: evidence from firearm background checks. *Southern Economic Journal*, 83(1):45–68, 2016.

John R Lott, Jr and David B Mustard. Crime, deterrence, and right-to-carry concealed handguns. *The Journal of Legal Studies*, 26(1):1–68, 1997.

Michael Luca, Deepak Malhotra, and Christopher Poliquin. The impact of mass shootings on gun policy. *Journal of public economics*, 181:104083, 2020.

Topher L McDougal, Daniel Montolio, and Jurgen Brauer. Modeling the us firearms market: the effects of civilian stocks, legislation, and crime. *International social science journal*.

Topher L McDougal, Daniel Montolio, Jurgen Brauer, et al. *Modeling the US firearms market: the effects of civilian stocks, crime, legislation, and armed conflict*. Institut d'Economia de Barcelona, Facultat d'Economia i Empresa, Universitat . . . , 2020.

Matthew Miller, Wilson Zhang, and Deborah Azrael. Firearm purchasing during the covid-19 pandemic: results from the 2021 national firearms survey. *Annals of Internal Medicine*, 175(2):219–225, 2022.

Sarah Moshary, Bradley Shapiro, and Sara Drango. Preferences for firearms and their implications for regulation. *American Economic Review: Insights*, forthcoming.

Aviv Nevo. Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics*, pages 395–421, 2000.

Martin O'Connell and Kate Smith. Optimal sin taxation and market power. Technical report, IFS Working Paper, 2021.

Thomas R Oliver. The politics of public health policy. *Annu. Rev. Public Health*, 27:195–233, 2006.

Kim Parker, Juliana Menasce Horowitz, Ruth Igielnik, J Baxter Oliphant, and Anna Brown. The demographics of gun ownership. *Pew Research Center*, 22, 2017.

Cora Peterson, Gabrielle F Miller, Sarah Beth L Barnett, and Curtis Florence. Economic cost of injury—united states, 2019. *Morbidity and Mortality Weekly Report*, 70(48):1655, 2021.

Rajeev Ramchand and Jessica Saunders. The effects of the 1996 national firearms agreement in australia on suicide, homicide, and mass shootings. *Contemporary Issues in Gun Policy: Essays from the RAND Gun Policy in America Project, Santa Monica, Calif.: RAND Corporation, RR-A243-2*, pages 43–65, 2021.

RAND. The relationship between firearm prevalence and violent crime. Technical report, 2018. URL `https://www.rand.org/research/gun-policy/analysis/essays/firearm-prevalence-violent-crime.html`.

Adam Rosenberg. The relationship between firearm prevalence and violent crime. Technical report, 2023.

Rosanna Smart. Firearm and ammunition taxes. Technical report, 2018. URL `https://www.rand.org/research/gun-policy/analysis/essays/firearm-and-ammunition-taxes.html#fnb7`.

Rosanna Smart, Andrew R. Morral, and Terry L. Schell. *The Magnitude and Sources of Disagreement Among Gun Policy Experts: Second Edition*. RAND Corporation, Santa Monica, CA, 2021. doi: 10.7249/RRA243-3.

Rebeccah L Sokol, Marc A Zimmerman, Laney Rupp, Justin E Heinze, Rebecca M Cunningham, and Patrick M Carter. Firearm purchasing during the beginning of the covid-19 pandemic in households with teens: a national study. *Journal of behavioral medicine*, 44: 874–882, 2021.

David M Studdert, Yifan Zhang, Erin E Holsinger, Lea Prince, Alexander F Holsinger, Jonathan A Rodden, Garen J Wintemute, and Matthew Miller. Homicide deaths among adult cohabitants of handgun owners in california, 2004 to 2016: a cohort study. *Annals of Internal Medicine*, 2022.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.

Kenneth E Train. *Discrete choice methods with simulation*. Cambridge university press, 2009.

Voting and Election Science Team. 2016 Precinct-Level Election Results, 2018. URL `https://doi.org/10.7910/DVN/NH5S2I`.

Douglas J Wiebe, Robert T Krafty, Christopher S Koper, Michael L Nance, Michael R Elliott, and Charles C Branas. Homicide and geographic access to gun dealers in the united states. *BMC Public Health*, 9:1–10, 2009.

	(1) Handguns	(2) Longguns
Caliber (Inches)	0.363	0.453
	(0.093)	(0.217)
Barrel Length (inches)	4.452	22.585
	(2.169)	(4.530)
High Capacity Weapon	0.188	0.154
	(0.391)	(0.361)
Shotgun	0.000	0.396
	(0.000)	(0.489)
100% Grading Used Price ($)	905.741	1354.701
	(1099.337)	(2487.826)
MSRP ($)	958.340	1635.588
	(663.283)	(3033.530)
Number of Gun Models	2835	4781

Table 1: Gun Characteristics summary statistics

	(1) Adult Pop.	(2) Potential Gun Buyers	(3) Gun Buyers	(4) Handgun Buyers	(5) Longgun Buyers
Female	0.513	0.225	0.085	0.107	0.047
Fraction White in Zipcode	0.681	0.687	0.813	0.804	0.828
Poverty Rate in Zipcode	0.105	0.103	0.081	0.083	0.078
Fraction BA+ in Zipcode	0.410	0.402	0.402	0.399	0.406
Density of Zipcode	1,437	1,236	665	710	587
Median Zipcode Income	79,878	78,950	88,456	87,961	89,304

Table 2: Demographics of Gun Owners in Massachusetts

	(1) IHS(Purchases)	(2) IHS(Purchases)	(3) IHS(Purchases)
Log(MSRP)	-0.579	-2.265**	-2.386**
	(0.446)	(1.032)	(1.038)
Observations	4465	4465	4465
R-Squared	0.73	-0.00	-0.00
Gun Model FE	Yes	Yes	Yes
Year FE	Year	Year	Class x Year
# Selected Instruments		18	17
Sup-Score Test Statistic		14.48	14.60
First Stage Sig. (p-value)		0.0000	0.0000
Method	OLS	IVLASSO	IVLASSO

Table 3: Estimated Price Elasticity of Demand for Firearms

	Used vs New Good Nesting	Class-level Nesting
Handgun	0.238	0.291
	(0.044)	(0.090)
	[0.1130]	[0.0000]
Longgun	0.325	0.365
	(0.080)	(0.255)
	[0.3096]	[0.0185]
Inside Good		0.896
		(0.098)
		[0.1445]

Table displays estimates of $\rho_{c,u}$ under Used vs New Good Nesting, and Estimates of ρ_c for Handguns and long guns under Class-level Nesting, while ρ_0 estimates are displayed in the row corresponding to inside good. Estimates are displayed, followed by standard errors in parentheses underneath. In brackets we display p-values for the following tests: for ρ_0, the test that $\rho_0 \leq 1$, for ρ_c, the test that $\rho_c \leq \rho_0$, and for ρ_c^u, the test that $\rho_c^u \leq \rho_c$. These significance tests are performed using a one-sided z-score test, treating the upper bound of the test as a known number.

Table 4: Nesting Parameters

Characteristic	Weapon Class	All	Frac Conservative	Female	Frac BA+ Edu	Frac White	Log(Median Inc)	Log(Density)
Constant	Handgun	3,014.04	1,796.53	-3,340.72	-2,713.21	2,611.14	367.93	-78.08
	Longgun	4,564.70	839.28	-4,653.42	-2,513.37	4,640.81	91.57	-106.64
Barrel Length (Inches)	Handgun	-278.72	-85.23	-200.70	45.37	49.70	-10.15	-4.13
	Longgun	-101.84	-13.78	-16.34	43.59	92.58	-9.09	-1.95
Caliber (Inches)	Handgun	-1,292.18	486.16	-1,321.93	-764.82	-685.08	79.24	15.36
	Longgun	-1,722.73	67.84	-265.33	-395.28	-608.95	-136.16	10.52
High Capacity Weapon	Handgun	410.03	37.29	-152.02	148.50	-275.22	129.11	8.81
	Longgun	-454.90	64.47	-191.05	250.01	-391.27	300.09	17.45
New Gun	Handgun	-1,946.06	251.30	1,288.06	-47.15	1.33	-8.55	11.86
	Longgun	-1,977.11	277.75	425.47	-1,026.18	-1,578.31	163.71	28.40
Shotgun	Longgun	883.89	-172.36	201.53	269.91	-12.60	128.07	7.42

Table displays estimates of consumer preferences over firearm characteristics, in terms of 2022 dollars. For the All column, the baseline taste for a preference, we display β_c/α, and for columns referring to a particular demographic, we display Π_c/α.

Table 5: Preferences for Gun Characteristic (in \$)

Panel A: Intent-To-Treat			
	(1)	(2)	(3)
	Incidents	Deaths	Injuries
Log(1+# Gun Shops)	0.061***	0.042**	0.071**
	(0.023)	(0.018)	(0.028)
Observations	165573	140040	130302
R-Squared	Yes	Yes	Yes
ZCTA5 Fixed Effects	Yes	Yes	Yes
State-Year Fixed Effects	Yes	Yes	Yes
Panel B: Non-linear Two-Sample IV			
	(1)	(2)	(3)
Log(1+# Transactions)	0.881	0.606	1.016
	[0.382, 2.220]	[0.165, 1.750]	[0.346, 2.473]
Observations (First Stage)	3759	3759	3759
Observations (Second Stage)	165573	140040	130302
First-Stage F Statistic	5.22	5.22	5.22
Bootstrapped Pr($\theta = 0$)	0.000	0.001	0.002
ZCTA5 Fixed Effects	Yes	Yes	Yes
State-Year Fixed Effects	Yes	Yes	Yes

Table 6: Gun Transactions and Gun Violence Incidents: Two-Sample IV Estimates

Note: Table presents two sample nonlinear IV estimates of the effect of firearm transactions on gun violence using the number of gun retailers as an instrumental variable. Column 1 is the count of incidents of gun violence. Column 2 is non-fatal gun violence injuries. Column 3 is gun assault fatalities. Estimates are computed on a balanced panel of zip code-year observations. Panel A displays estimates of the intent-to-treat effect from a poisson regression and Panel B displays estimates of the two-sample IV effect of firearm transactions on gun-related incidents using the number of gun retailers as an instrumental variable. Standard errors in parentheses in Panel A clustered by zip code. Panel B estimates are constructed via a first stage regression of retailers on transactions, followed by second stage poisson regression of predicted transactions across the U.S. 95% confidence intervals included in brackets in Panel B derive from a bias-corrected bootstrap procedure of 1000 replications, where resampling is done at the zipcode level and stratified by state. P-values are computed from the bias-corrected bootstrap distribution. * denotes $p < .1$, ** denotes $p < .05$, *** denotes $p < .01$, All regressions include fixed effects for zip code and state-year.

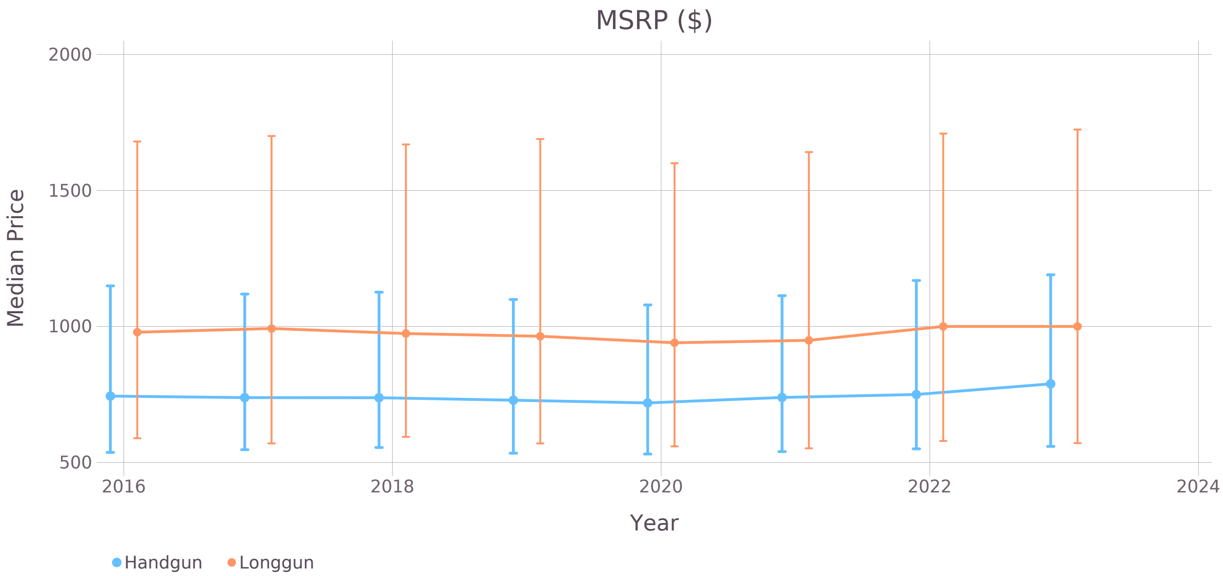

Figure 1: New Gun Prices over Time

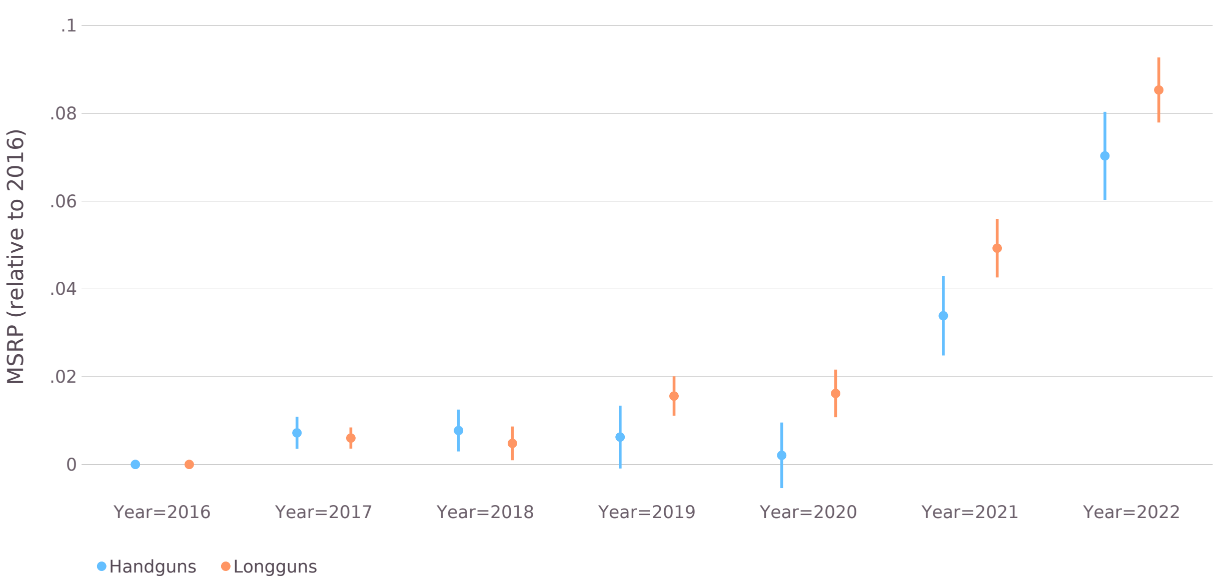

Figure 2: New Gun Prices over Time (Controlling for Variety)

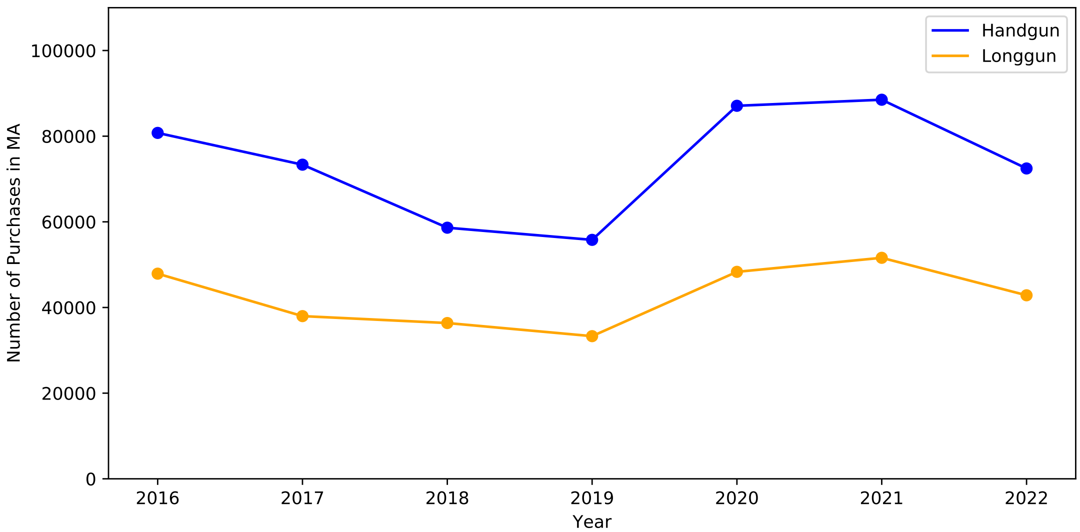

Figure 3: Aggregate Purchases by Weapon Class in Massachusetts

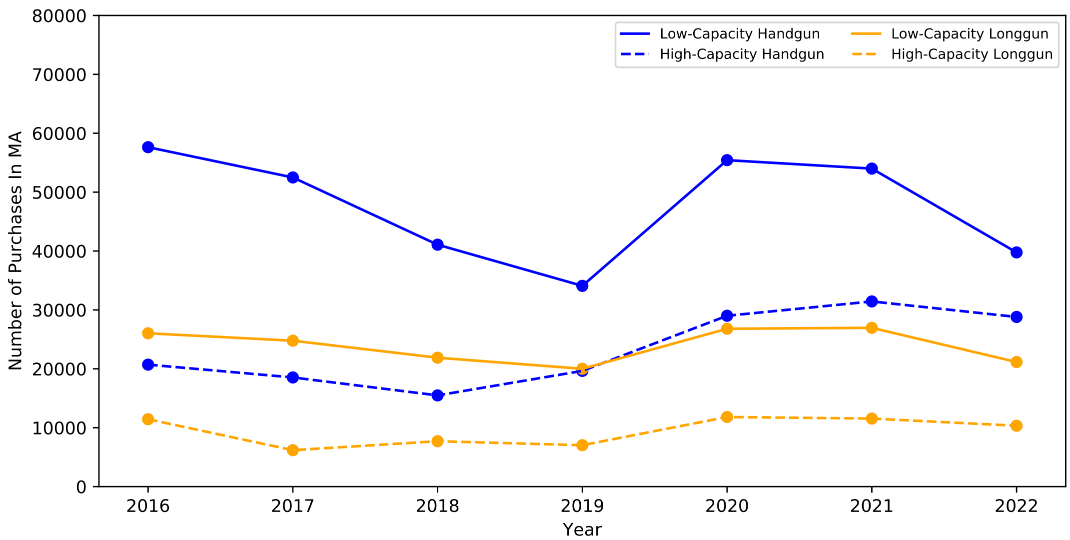

Figure 4: Aggregate Purchases by Weapon Class and High-Capacity Status in Massachusetts

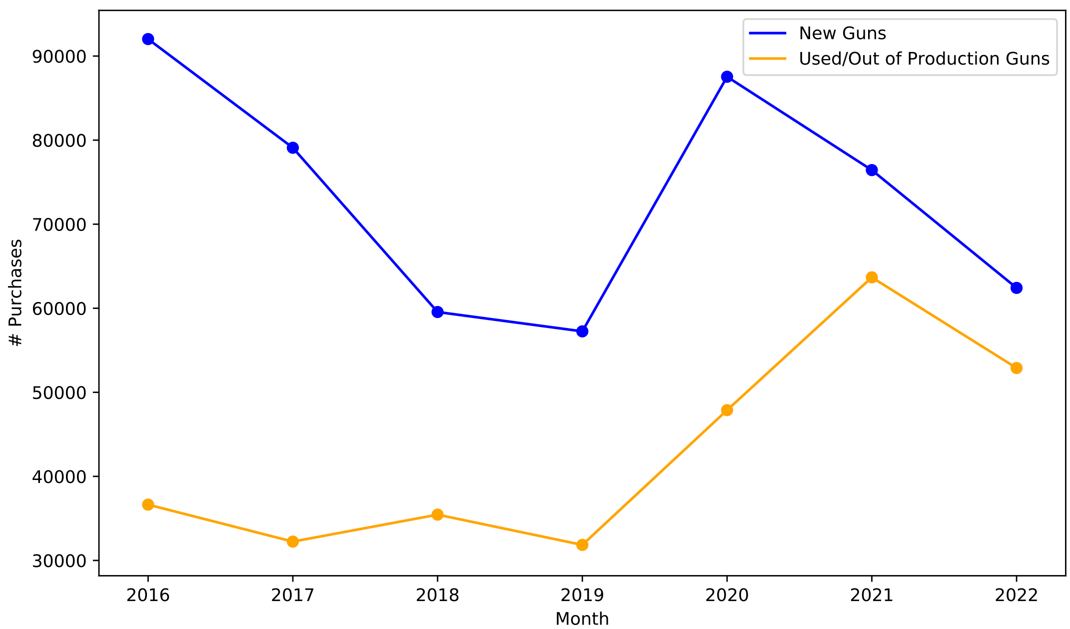

Figure 5: Aggregate Gun Purchases by Used vs New in Massachusetts

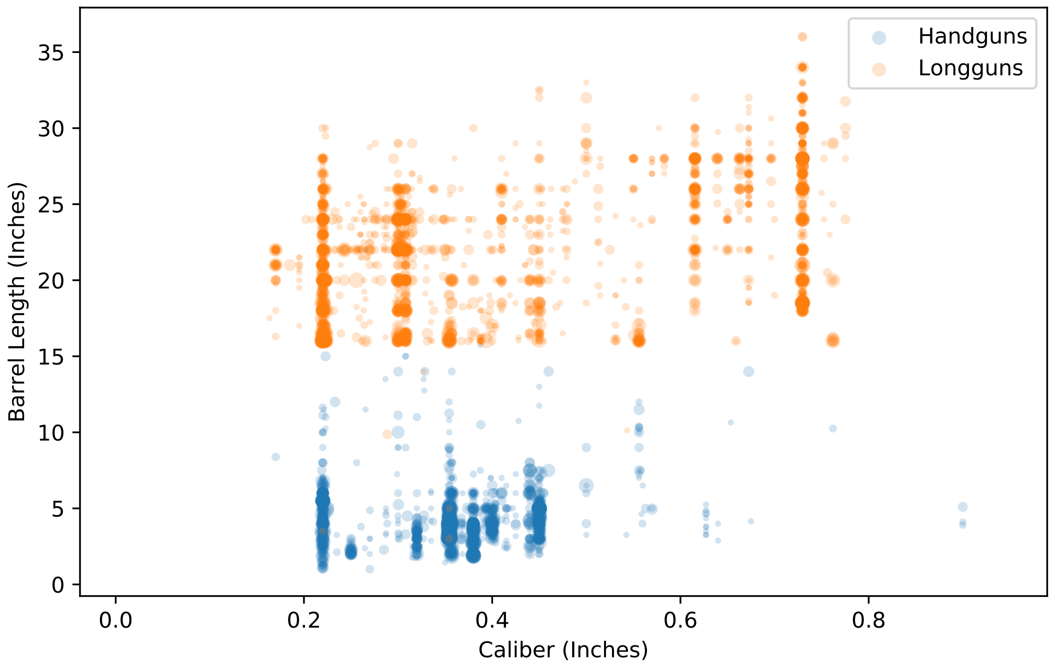

Figure 6: Caliber and Barrel Length Distribution Across Gun models

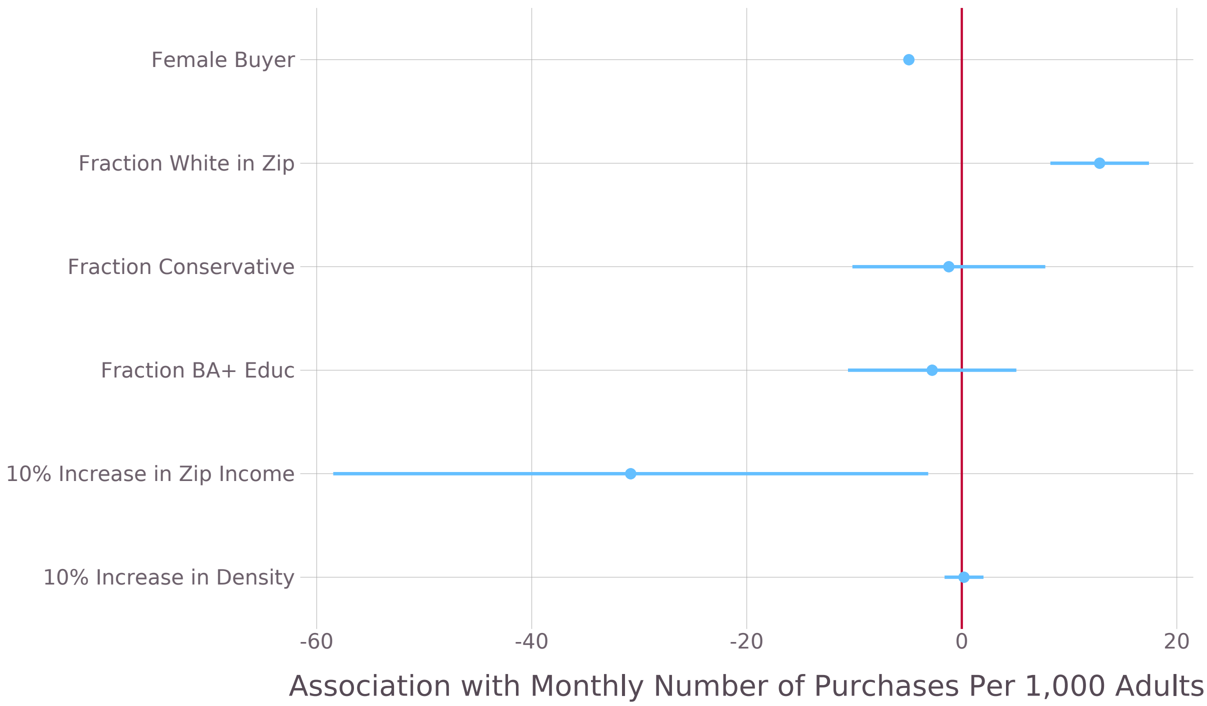

(a) Gun Purchase Frequencies (Extensive Margin)

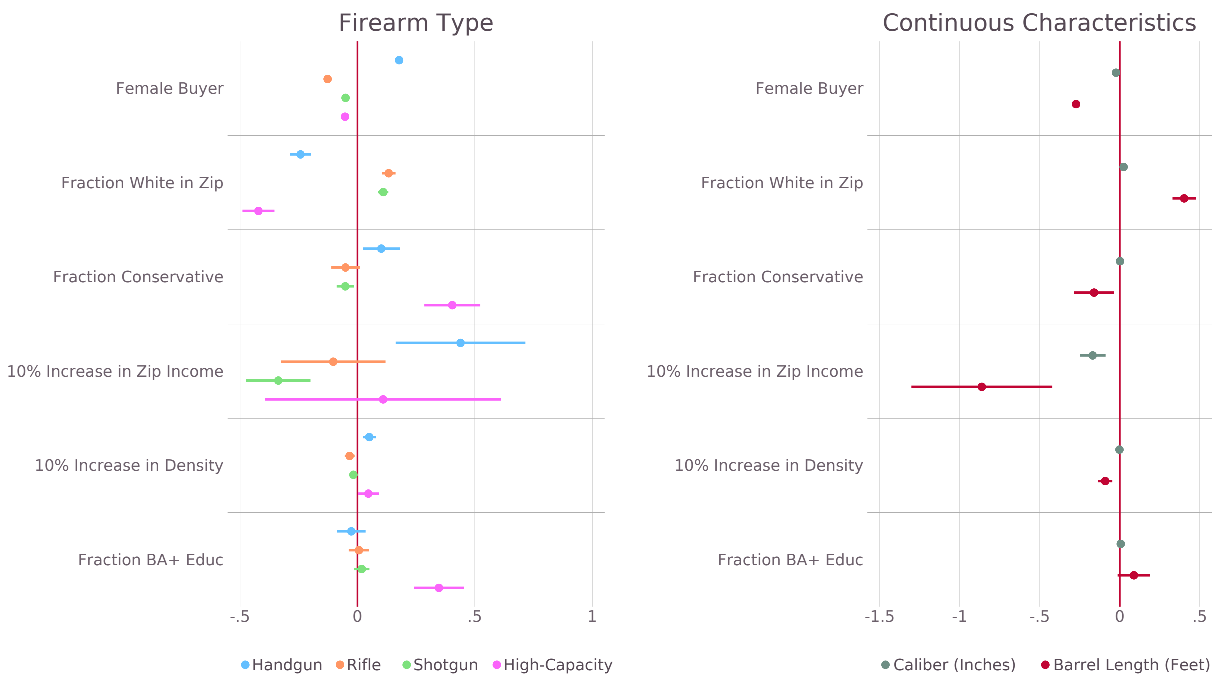

(b) Gun Characteristics (Intensive Margin)

Figure 7: Correlates of Demographics with Gun Purchases, Controlling for Market Structure

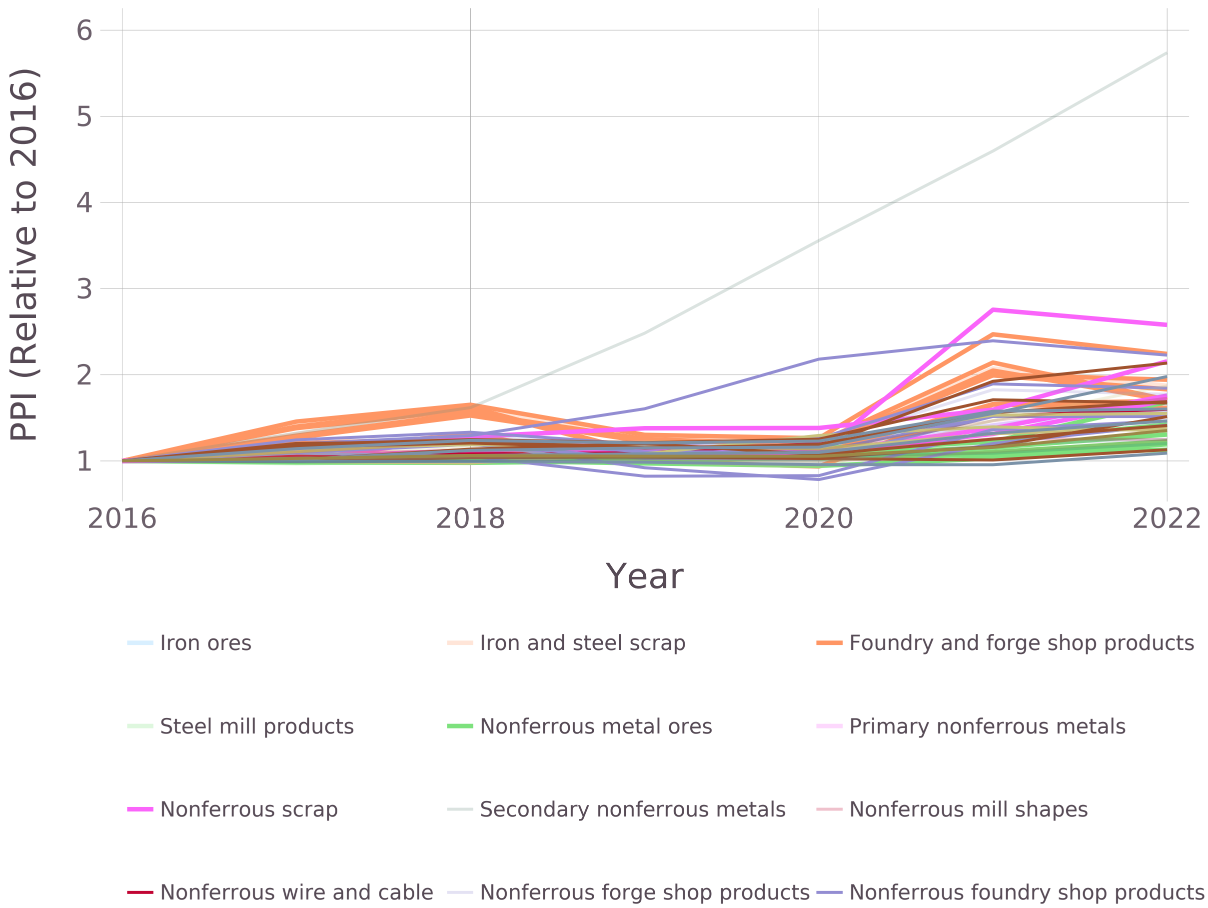

Figure 8: Time Series of Annual Price Indices for Firearms Production Inputs

(a) Estimated Non-Zero coefficients for IVLASSO of Prices

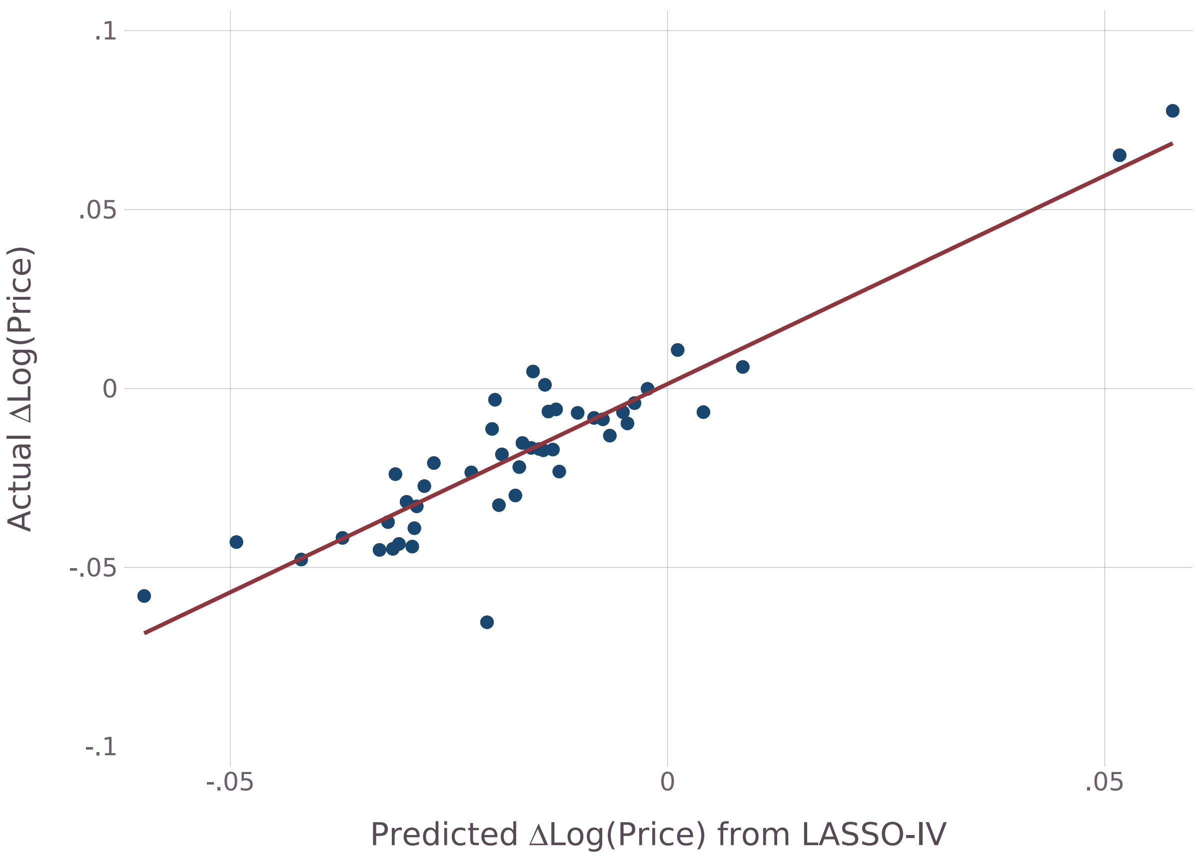

(b) Binscatter of Predicted IVLASSO and Actual Prices

Note: Panel (a) plots the magnitude of the selected coefficients from a lasso of log(MSRP) on logged commodity prices interacted with an ID for manufacturer, with gun model and weapon class times year fixed effects. Panel (b) plots the binscatter of year-to-year changes in residualized predicted and actual MSRP, for gun models produced by manufacturers with non-zero coefficients.

Figure 9: First Stage Estimates for IVLASSO of Prices

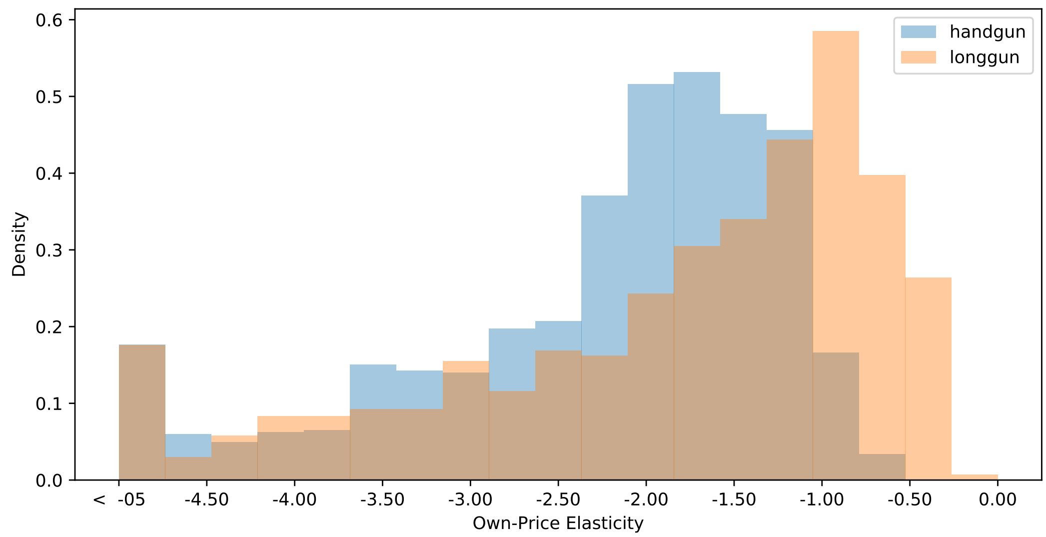

(a) Distribution of Own-Price Elasticities in Massachusetts, by Gun Class

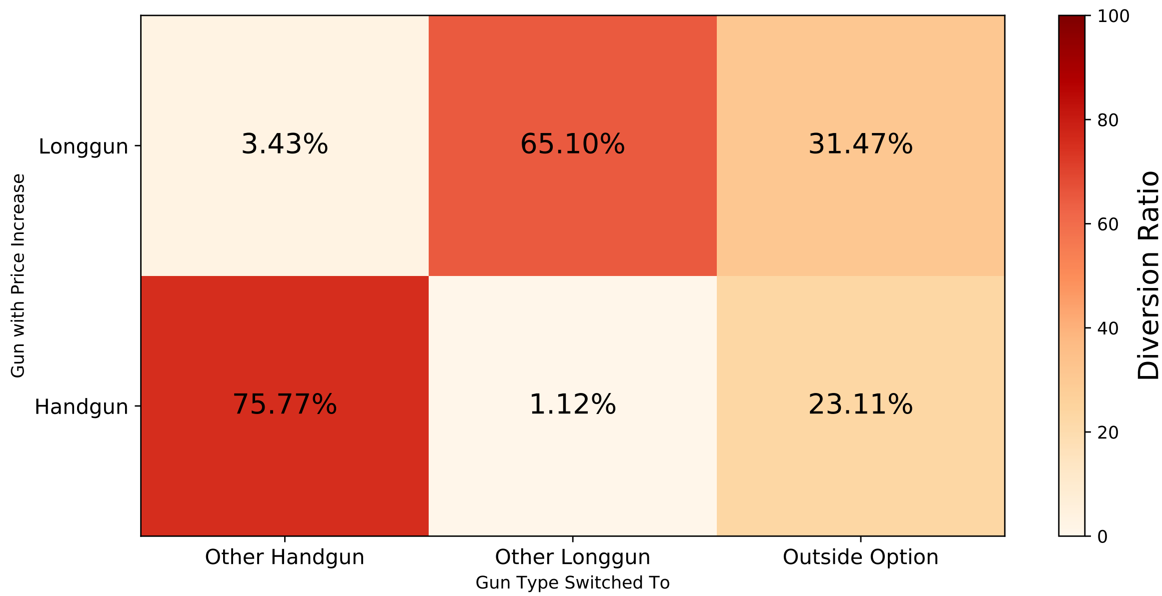

(b) Substitution Patterns (Diversion Ratios) across Gun Classes in Massachusetts

Figure 10: Own and Cross Firearm Substitution Patterns

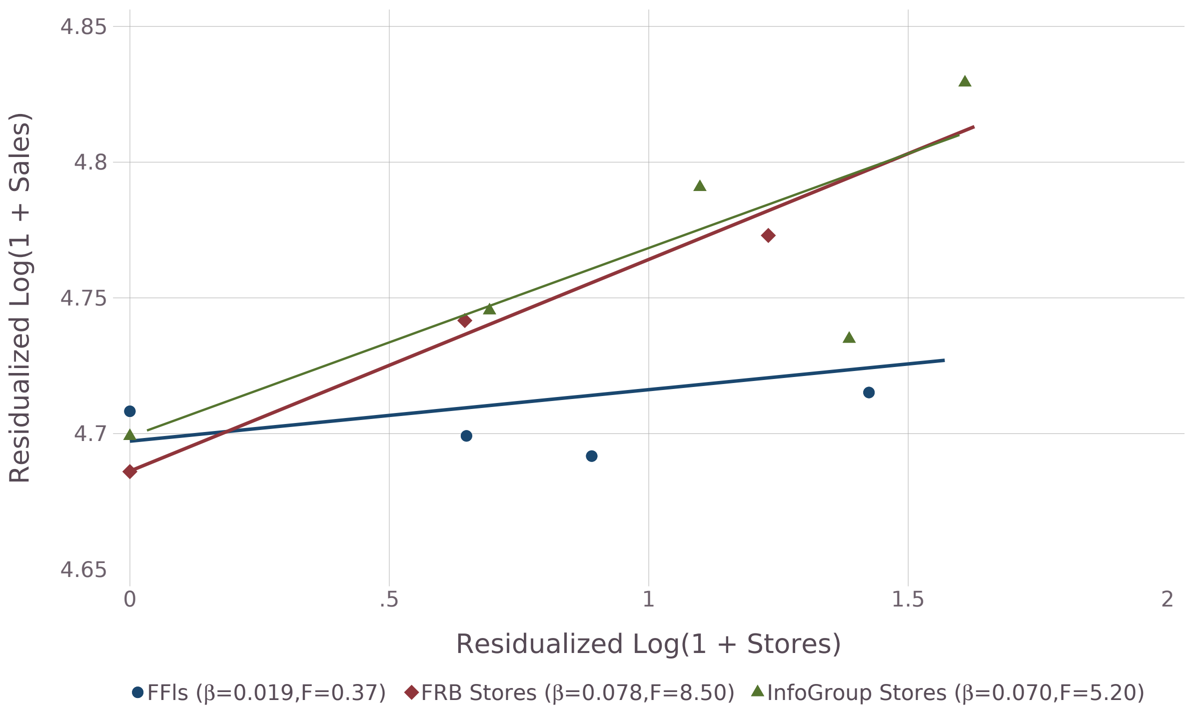

Figure 11: Gun Retailer Entry and Consumer Firearm Transactions

Note: Figure plots residualized binscatter between firearm retailers and firearm transactions within a zip code-year in pre pandemic data. Residualization is by zip code and year, with binning selected by the data driven procedure of Cattaneo et al. [2019]. Coefficients F-statistics, and best-fit lines are from a regression on the underlying balanced panel of zip code-year observations and include fixed effects for 5-digit zip code and period. Blue circles corresponds to measures of firearm retailers using the count of federal firearms licensees. Red diamonds corresponds to measures of firearm retailers using the active firearm retailers from the transaction data in Massachusetts. Green triangles correspond to measures of firearm retailers using specialty firearms data from Data Axle (formerly InfoGroup).

Panel A: Massachusetts

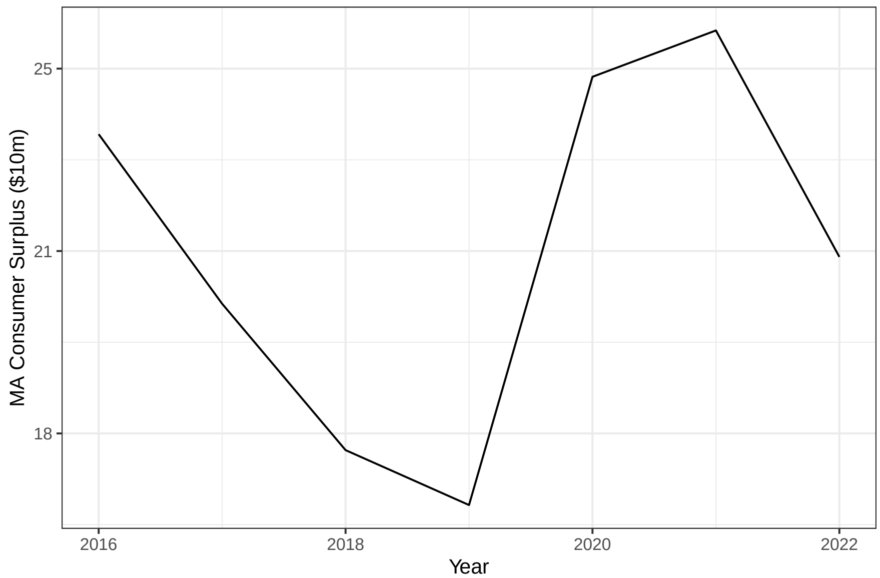

Panel B: U.S.

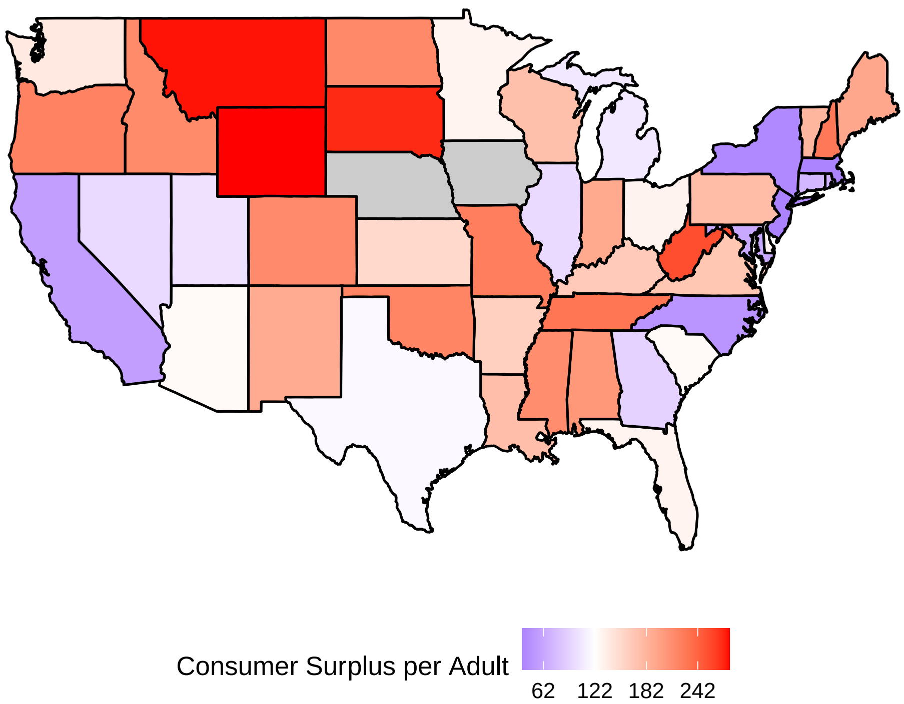

Figure 12: Consumer Surplus from the Firearm market in Massachusetts and the U.S.

Figure shows dollar value of consumer surplus from the firearms market. Panel A shows overall consumer surplus in Massachusetts each year. Panel B shows the dollar value of consumer surplus per adult in each state, averaged over the periods in Panel A. White represents the population-weighted mean consumer surplus per adult in the U.S. Scale ticks represent population-weighted standard deviations of consumer surplus per adult.

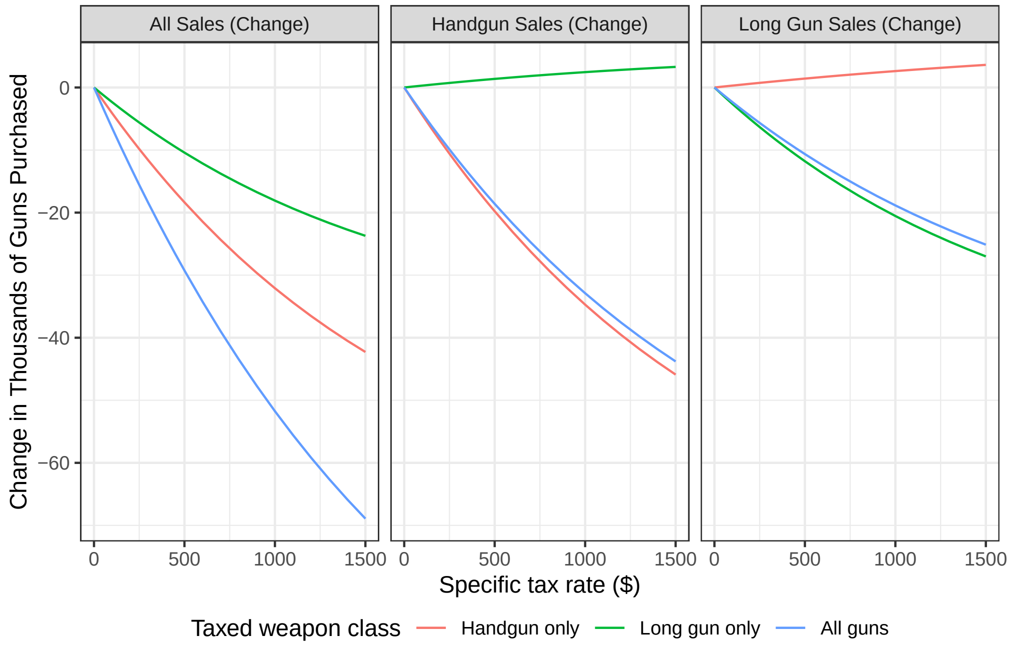

(a) Firearm Sales

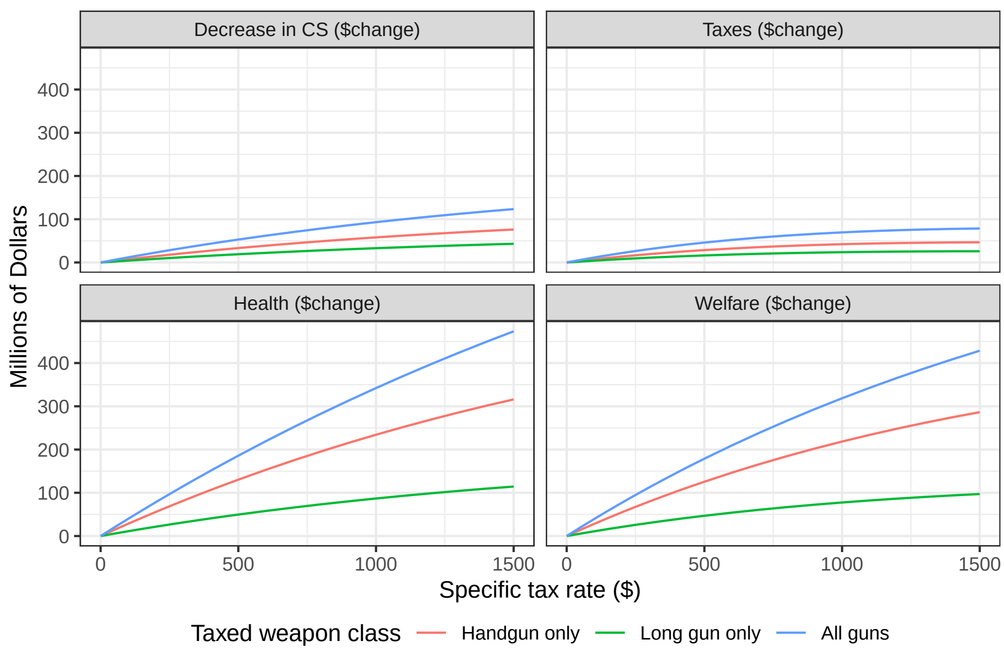

(b) Welfare Components

Figure 13: Intermediate Outcomes as a Function of Taxes on Handguns and Long guns in Massachusetts

Figure shows changes in intermediate outcomes as a function of the specific tax rate on handguns and long guns in Massachusetts. Horizontal axis is the tax amount. Color represents the weapon class being taxed. Vertical axis is the value of the intermediate outcome. Intermediate outcomes in the top sub-figure are the change in total firearm sales, handgun sales, and long gun sales. In the bottom subfigure are dollar changes in tax revenue, consumer surplus, the value of gun violence avoided, and overall welfare.

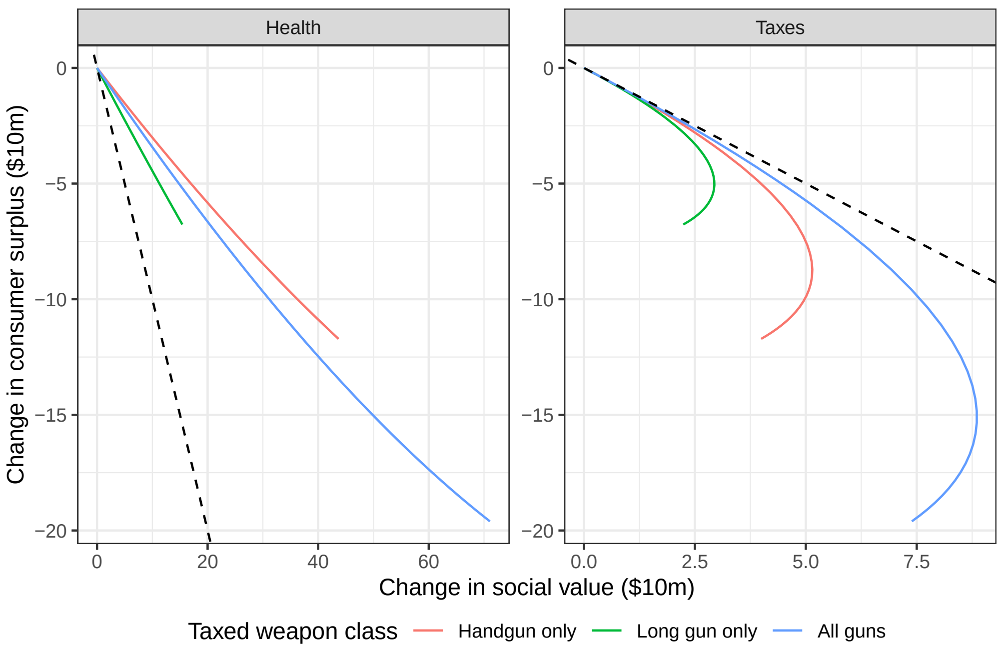

Figure 14: Policy Frontiers Generate by Taxes on Handguns and Long guns

Figure shows frontier between private-market consumer surplus and social value generated by alternative taxes on the Massachusetts firearm market. Vertical axis is the change in consumer surplus from the no-tax status quo. Horizontal axis is the change in social value from gun violence (left panel) and tax revue multiplied by a Marginal Value of Public Funds of 0.15 (right panel). Color is the tax rate on long guns. Dotted line is negative 45-degree line.

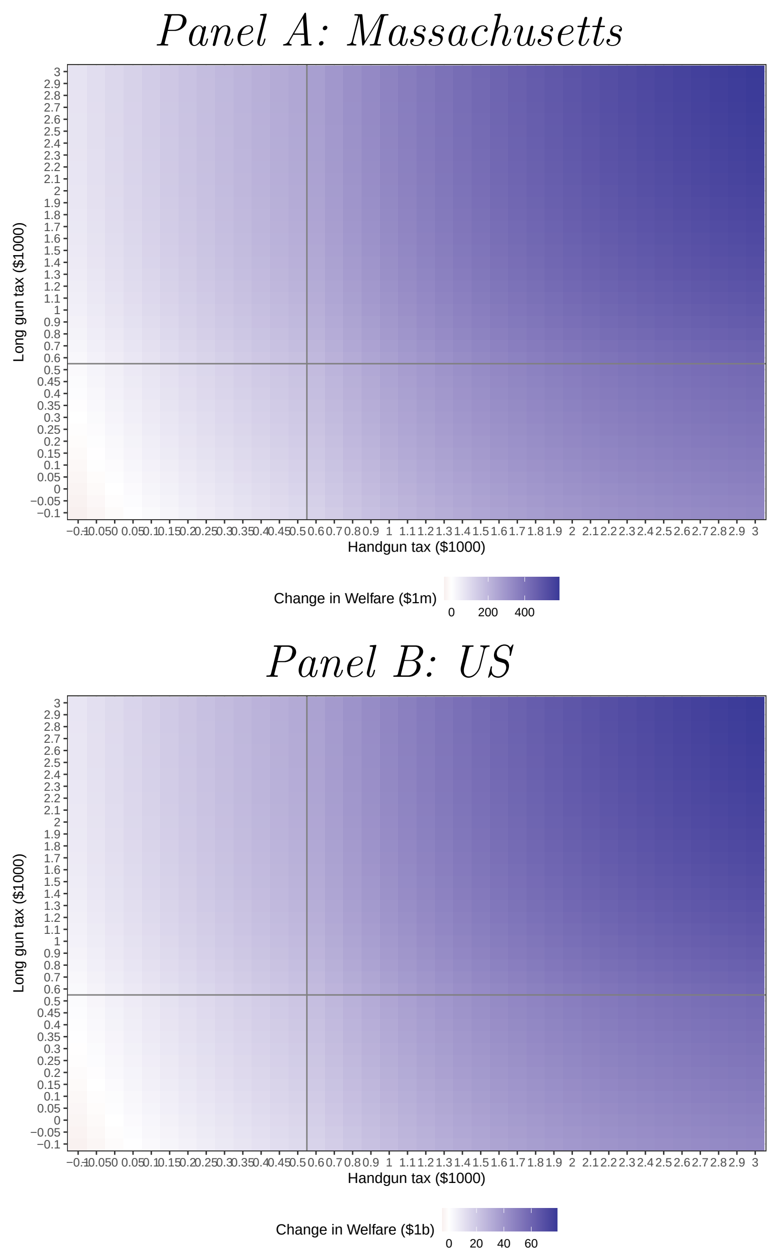

Figure 15: Changes in Welfare as a Function of Taxes on Handguns and Long Guns

Figure shows heat map of change in social welfare from counterfactual taxes on handguns and long guns. Panel A is outcomes in Massachusetts, Panel B is outcomes in the U.S. Horizontal axis is handgun tax. Vertical axis is long gun tax. Color represents the change in social welfare from the no-tax status quo. Gray lines denote a change in the scaling of their respective axes.

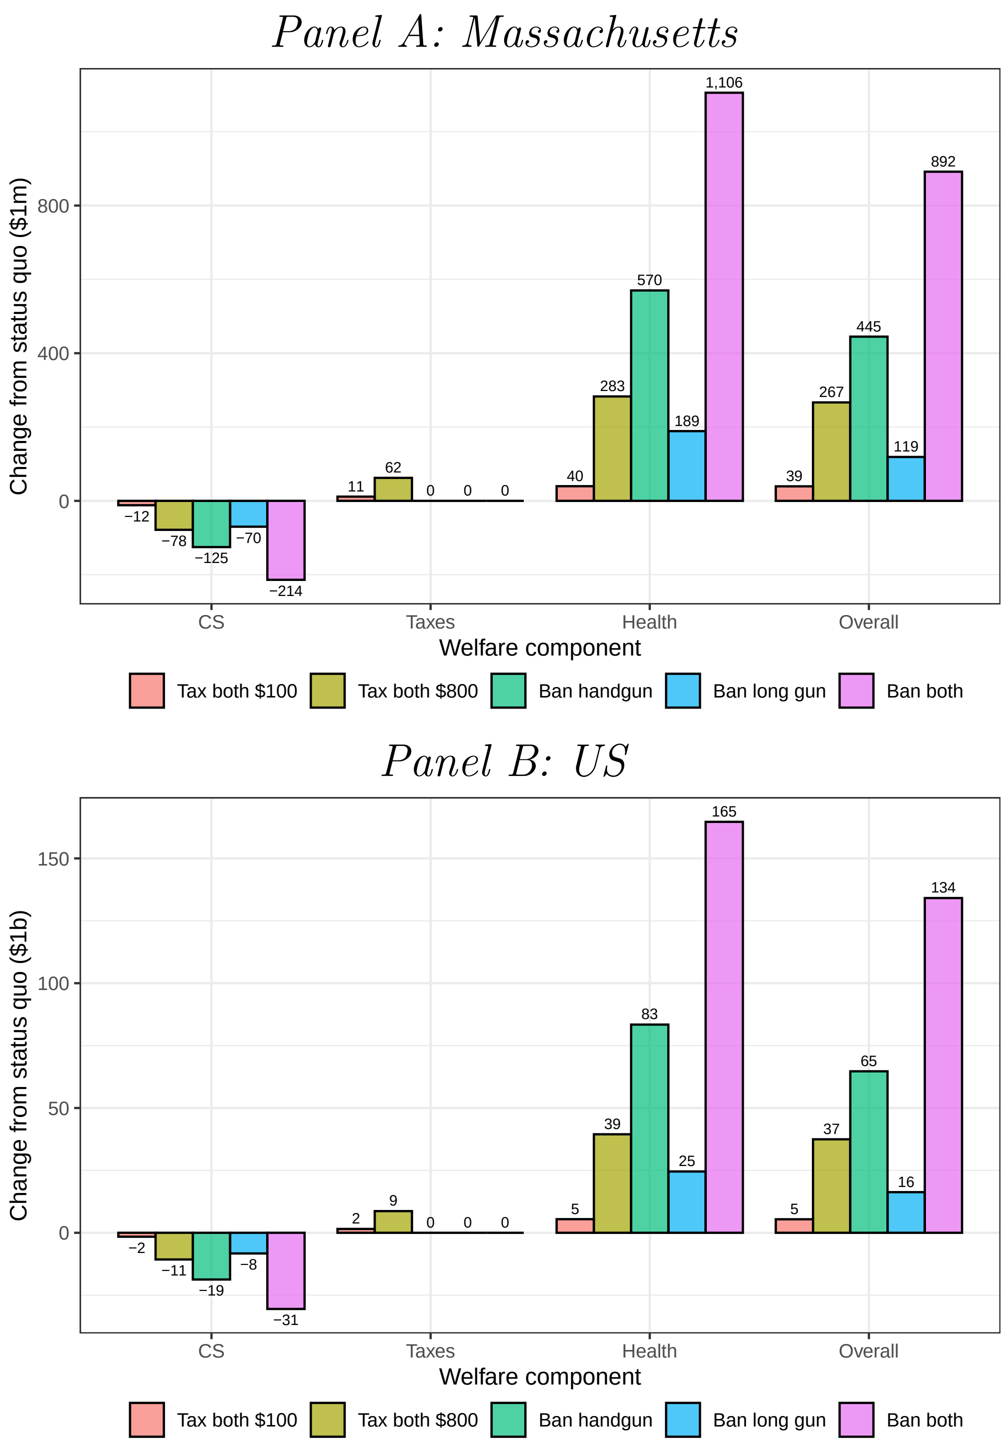

Figure 16: Changes in Welfare from Gun Bans

Figure shows change social welfare from counterfactual bans on handguns and long guns. Panel A is Massachusetts and Panel B is U.S. Color indexes counterfactual taxes or bans on handguns and long guns. Horizontal axis is welfare component. Vertical axis is its change from status quo in dollars of social value.

A Matching of Blue Book and FRB Gun Models

In this section, we describe our matching procedure to merge models in the FRB to BBGV gun IDs. First, we manually match all makes, by unique (lowercase) string, with at least 30 transactions in our time period in Massachusetts to their corresponding manufacturer ID in BBGV. 95% of all transactions in FRB are matched to a BBGV manufacturer ID.

We then preprocess gun model names in both datasets. We remove non-alphanumeric characters, remove the standalone word "model", which we view as uninformative, and appears in many FRB and BBGV model names, convert roman numerals to numbers (e.g. $II \rightarrow 2$), and then tokenize the string so that words appear in alphabetical order, followed by any tokens that are numbers. This latter point is done to ensure that the model numbers, which are often the key identifiers of guns, occupy the same part of the string and therefore the matching algorithm weights these corresponding more closely. We then match processed gun names, in the FRB, to the most similarly named gun in BBGV. The candidate BBGV models are constrained to be within manufacturer and weapon class, as well as have price data (new or used) in the years we observe transactions in the FRB.[45] We use the weighted ratio score of the `fuzzywuzzy` Python package to match strings, and retain matches with at least a 60% similarity score. We choose this relatively low threshold to ensure as many matches as possible, and we do not ignore some transactions in our demand estimation simply due to a low quality string match; though this will introduce additional estimation error to our demand system. We proceed assuming the matched BBGV gun model is the true gun model represented in the FRB data. If there are ties, we break them by an indicator of whether the gun is in active production (has an MSRP for the year), followed by the number of years for which MSRP data is available. In total, through this procedure, we are able to match 90% of FRB transactions to a BBGV ID, for a total of 7,616 unique BBGV gun models. About 70% of transactions occur in years where the BBGV models are actively produced.

Finally, we complete the merge by assigning each transaction a common set of gun characteristics by BBGV model ID, taken as the median FRB value in the FRB, across transactions. As a validation of the merge, the imputed values agree with the recorded gun model characteristics: the imputation for high-capacity agrees 80% of the time, and the standard deviation of the imputation error for barrel length is 85% lower than the standard deviation of barrel length.

[45] This latter condition is used to avoid matching FRB guns to BBGV models that were released after we observe transactions.

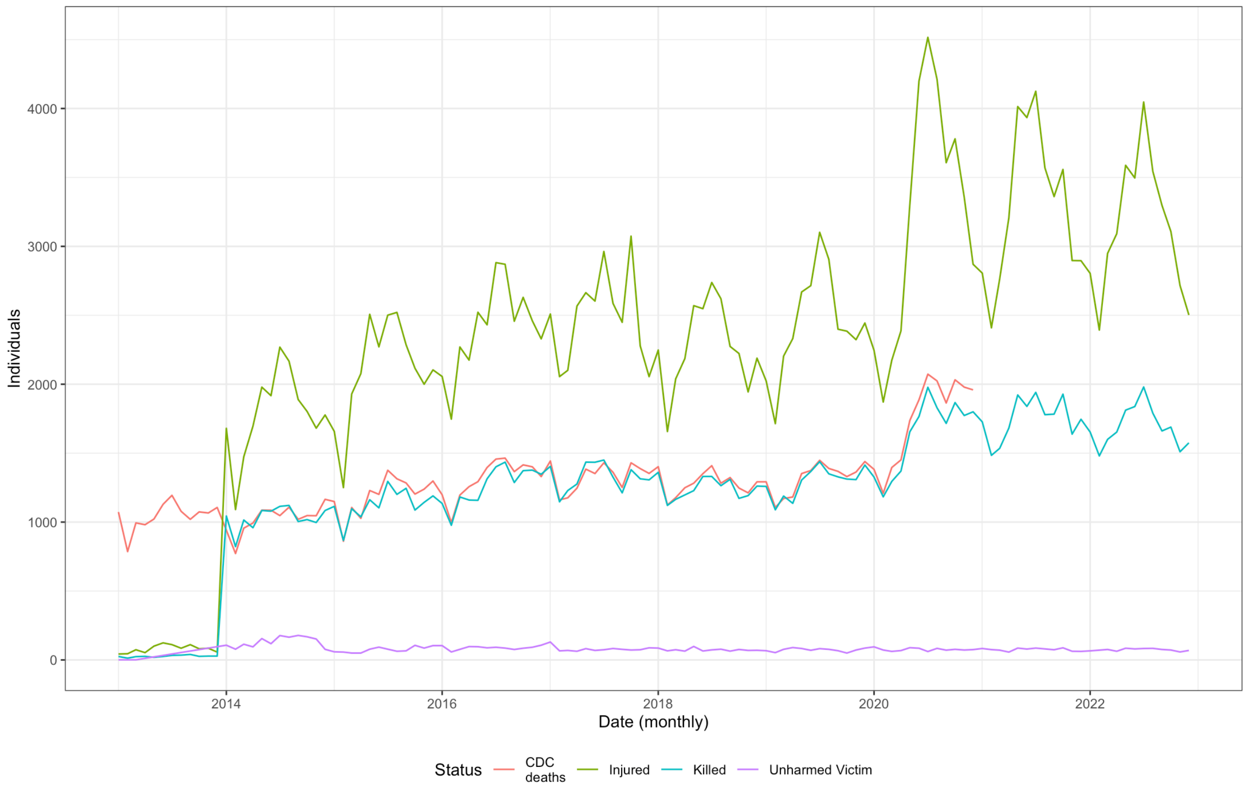

Figure A1: Incidents of Gun Violence in GVA and CDC Records
Figure shows time series of different measures of gun violence. Horizontal axis is the date in months. Vertical axis is the number of individuals affected by gun violence during a month. Color indexes measure of gun violence. Orange is deaths from gun-related assault fatalities in the U.S as recorded by the CDC under ICD10 codes X93, X94, and X95. Blue is the same quantity recorded by the GVA. Green is non-fatal injuries from gun violence in the U.S as recorded by the GVA. Purple is non-injurious incidents of gun violence in the U.S as recorded by the GVA.

	(1) Mean Util $\delta_{j,t}$	(2) Mean Util $\delta_{j,t}$
Price Coef. α	-0.00060***	-0.00059***
	(0.00017)	(0.00017)
Observations	2879	2879
R-Squared	-0.05	-0.05
Gun Model FE	Yes	Yes
Year FE	Year	Class x Year
# Selected Instruments	11	12
Sup-Score Test Statistic	11.12	11.20
First Stage Sig. (p-value)	0.0000	0.0000

Table A1: Estimated Price Coefficient From Structural Model

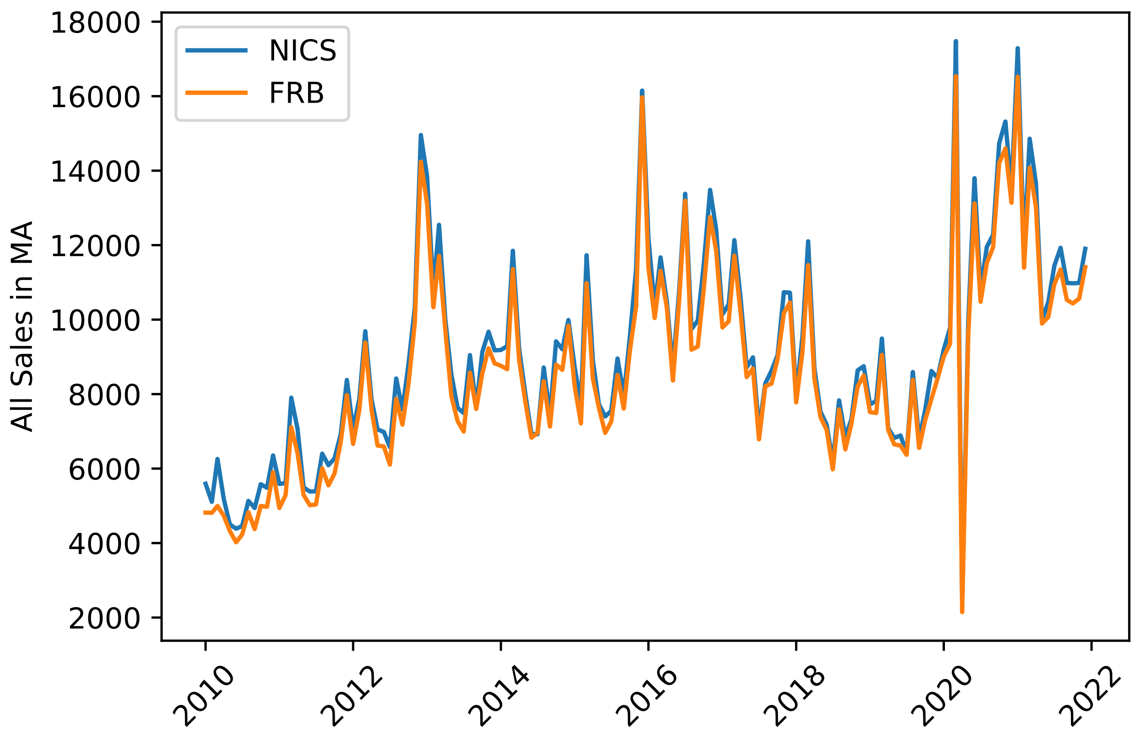

Figure A2: FRB Recorded Gun Transactions vs NICS Background Checks in Massachusetts, by Month

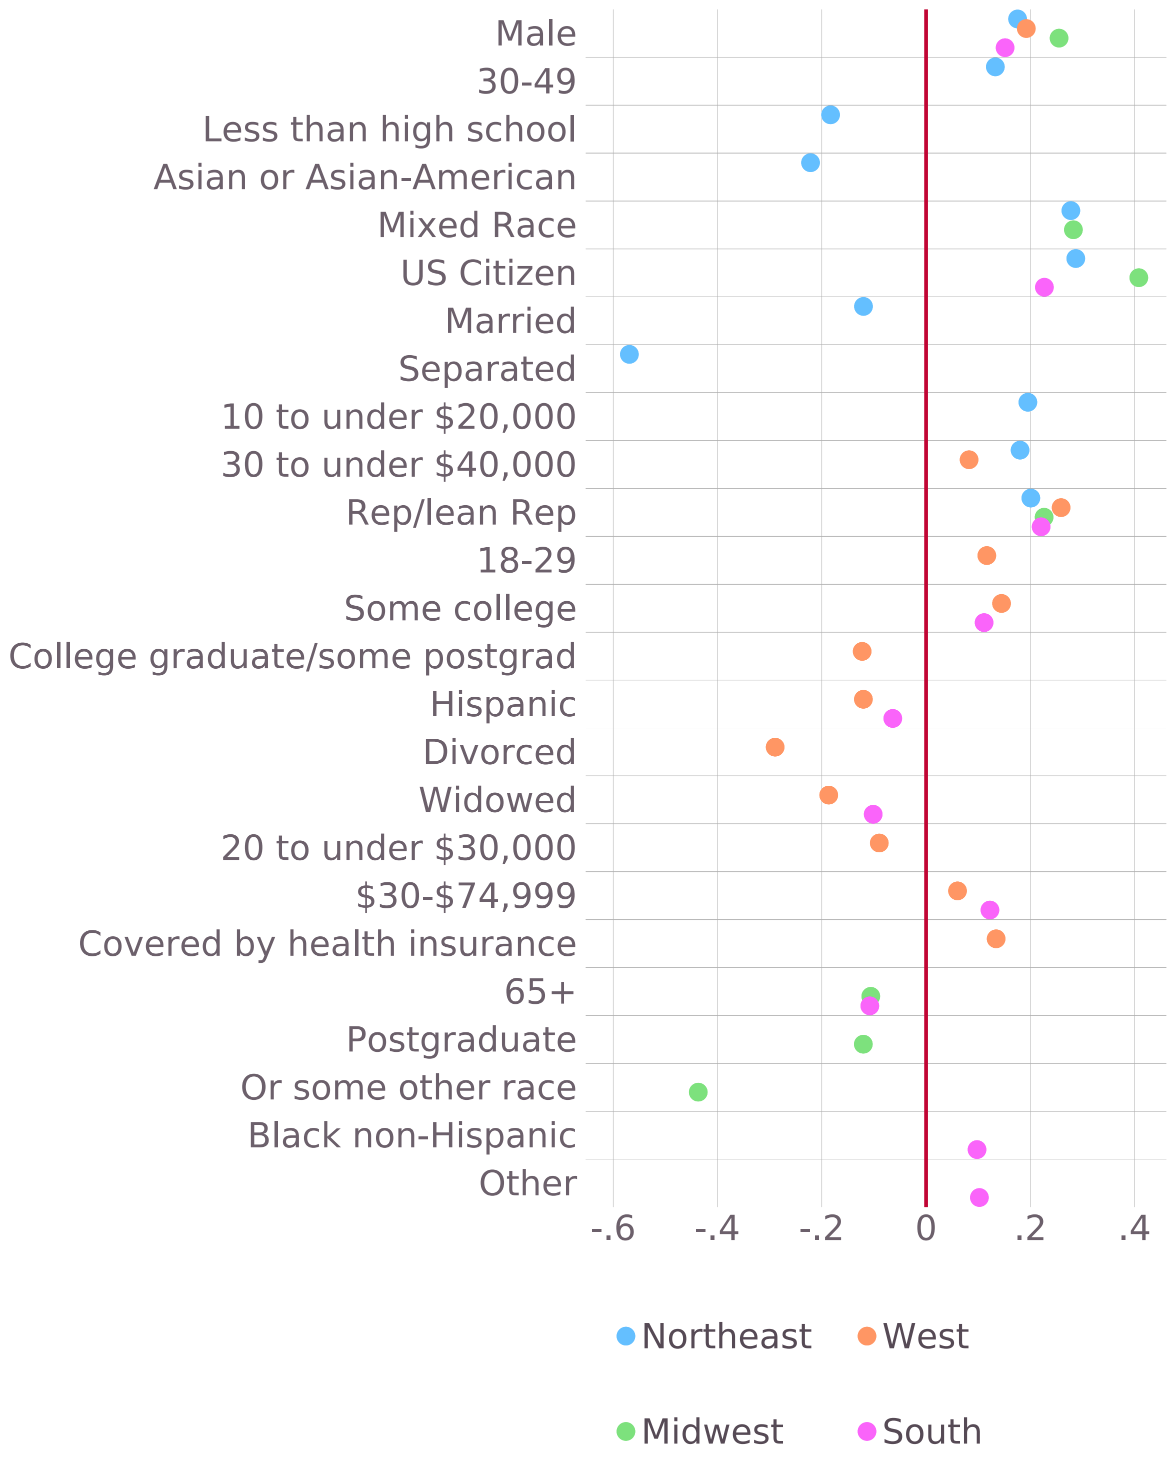

Figure A3: Estimated Lasso Coefficients for Predicting Market Size, From Pew's American Trends Panel

Characteristic	Weapon Class	All	Frac Conservative	Female	Frac BA+ Edu	Frac White	Log(Median Inc)	Log(Density)
Constant	Handgun	1.774***	1.057***	-1.966***	-1.597***	1.537***	0.217***	-0.046***
		(0.021)	(0.040)	(0.026)	(0.026)	(0.050)	(0.011)	(0.001)
	Longgun	2.686***	0.494***	-2.739***	-1.479***	2.731***	0.054***	-0.063***
		(0.022)	(0.050)	(0.062)	(0.045)	(0.091)	(0.018)	(0.002)
Barrel Length (Inches)	Handgun	-0.164***	-0.050***	-0.118***	0.027***	0.029***	-0.006*	-0.002***
		(0.001)	(0.012)	(0.022)	(0.009)	(0.007)	(0.003)	(0.001)
	Longgun	-0.060***	-0.008	-0.010***	0.026***	0.054***	-0.005**	-0.001***
		(0.000)	(0.006)	(0.003)	(0.008)	(0.014)	(0.002)	(0.000)
Caliber (Inches)	Handgun	-0.760***	0.286**	-0.778***	-0.450***	-0.403***	0.047	0.009*
		(0.018)	(0.138)	(0.145)	(0.144)	(0.104)	(0.050)	(0.005)
	Longgun	-1.014***	0.040	-0.156*	-0.233	-0.358**	-0.080	0.006
		(0.024)	(0.250)	(0.086)	(0.246)	(0.171)	(0.102)	(0.010)
High Capacity Weapon	Handgun	0.241***	0.022	-0.089***	0.087***	-0.162***	0.076***	0.005***
		(0.003)	(0.016)	(0.017)	(0.021)	(0.031)	(0.015)	(0.001)
	Longgun	-0.268***	0.038	-0.112***	0.147***	-0.230***	0.177***	0.010***
		(0.005)	(0.046)	(0.031)	(0.055)	(0.062)	(0.047)	(0.003)
New Gun	Handgun	-1.145***	0.148**	0.758***	-0.028	0.001	-0.005	0.007***
		(0.022)	(0.058)	(0.132)	(0.052)	(0.037)	(0.022)	(0.002)
	Longgun	-1.164***	0.163	0.250***	-0.604***	-0.929***	0.096*	0.017***
		(0.024)	(0.137)	(0.070)	(0.169)	(0.224)	(0.053)	(0.006)
Shotgun	Longgun	0.520***	-0.101	0.119**	0.159	-0.007	0.075	0.004
		(0.011)	(0.119)	(0.047)	(0.119)	(0.069)	(0.050)	(0.005)

Table displays estimates of consumer preferences over firearm characteristics, in terms of utils. For the All column, the baseline taste for a preference, we display β_c, and for columns referring to a particular demographic, we display Π_c. Below each parameter, standard errors are shown. * denotes $p < .1$, ** denotes $p < .05$, *** denotes $p < .01$, from a z-score test that the parameter is zero.

Table A2: Characteristic Preferences (in utils)

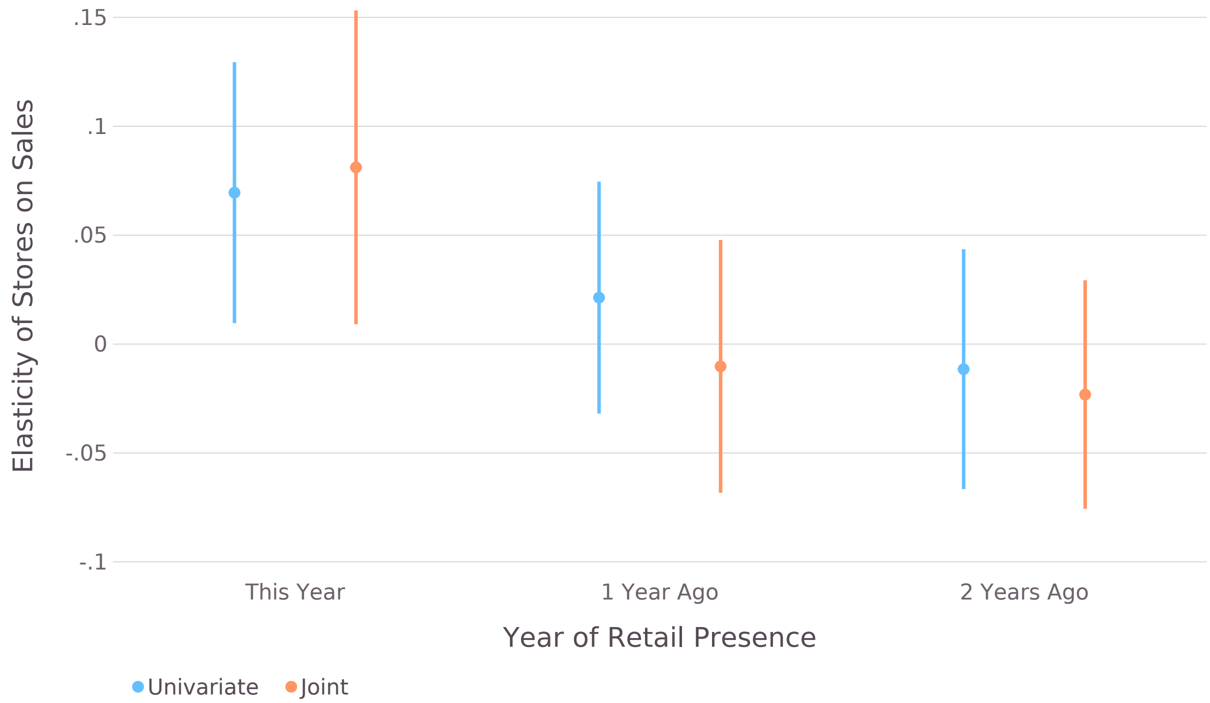

Figure shows coefficient estimates from a regression of gun transactions on gun retailers this year ($t = 0$), one year ago ($t = -1$), and two years ago $t = -2$. The underlying regression is a regression of $\log(1 + \text{Transactions})$ on $\log(1 + \text{Stores})$ with zipcode and year fixed effects. Univariate denotes regressions where only the single (lagged) retail presence is included, while joint denotes regressions where all three measures of retail presence are included in a single regression. Standard errors used to generate the displayed 95% confidence intervals are clustered at the zipcode level.

Figure A4: Effect of Current and Past Gun Retailers on Gun Transactions

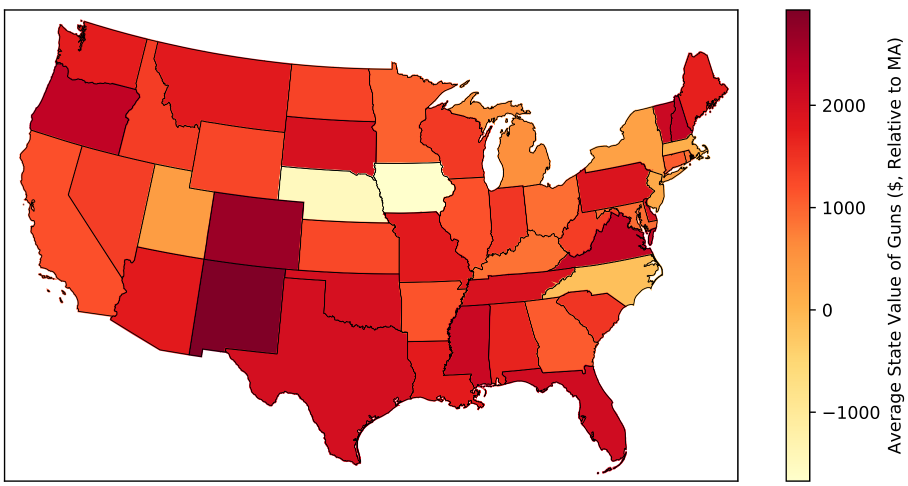

Figure shows the average value of τ_t, $\bar{\tau}_s = \sum_{t:state(t)=s} \tau_t$ for each U.S. state from 2016-2022, divided by the price coefficient α, to express the heterogeneity in willingness-to-pay for firearms across markets in dollar terms.

Figure A5: Average State-year Taste Shifters for Firearms, τ_t, by State, in 2022 $

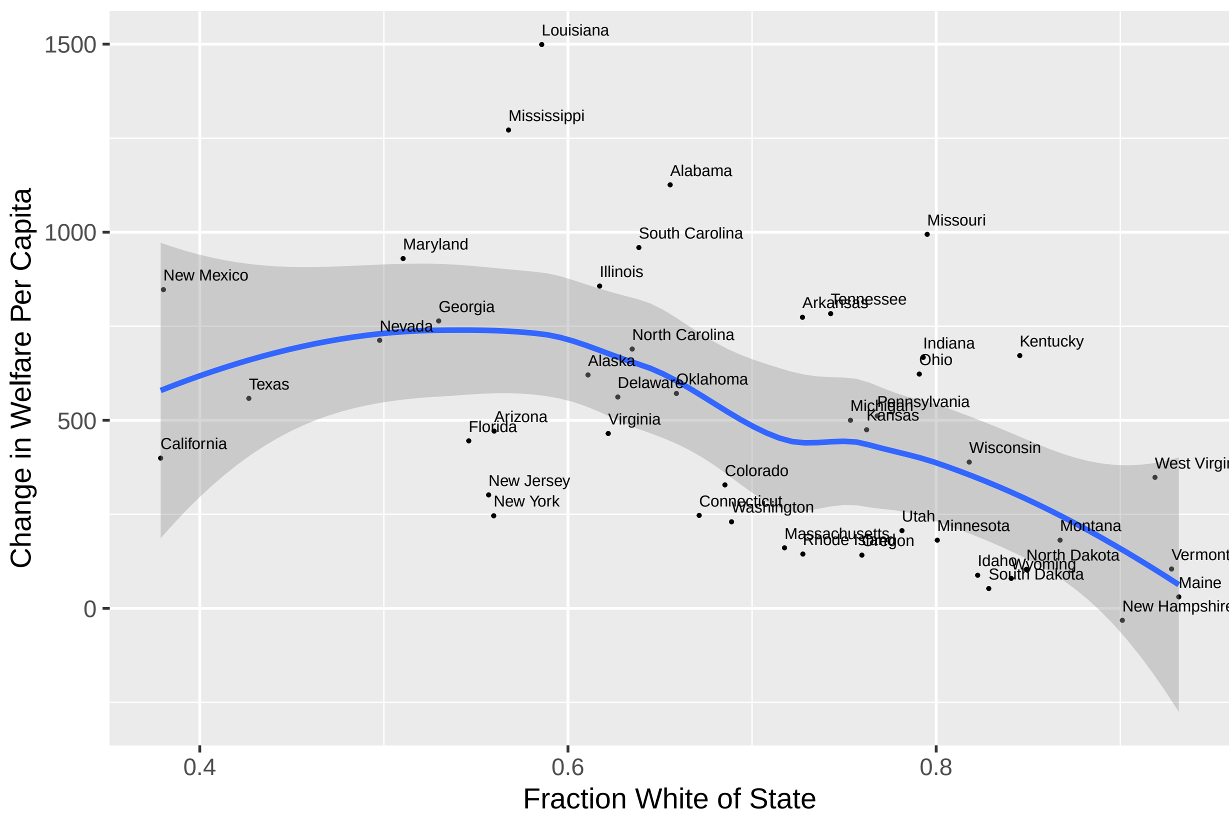

Figure shows the average change in total welfare per capita by state against the fraction of the adult population that is white.

Figure A6: Distribution of Welfare Changes from Ban on Firearm Sales, by Racial Composition