Second-Best Amendment: Market Power and Tax Design in the Firearms Industry*

Click Here for Latest Version

Luis Armona†, Adam M. Rosenberg‡

August 27, 2026

Abstract

This paper studies how market power and tax design determine surplus in the U.S. consumer firearms industry. Linking data on prices and quantities, we estimate an equilibrium model of the industry and its effect on homicide. Firm markups are substantial, but largely unrelated to firearms' homicide externalities. We then derive and simulate counterfactual firearm taxes under a "second-best" policy that respects the political influence of firearm consumers. This policy shifts taxes towards more harmful firearms, reducing homicides, while maintaining the aggregate surplus of market participants. It also redistributes surplus towards areas traditionally opposed to gun control.

JEL Codes: L13, L51, I18, H21, H23

Keywords: Firearms, Regulation, Externalities, Taxes, Diversion, Market Power, Political Economy

*We are grateful to the Massachusetts Firearms Records Bureau for their assistance in working with the FRB dataset, as well as Matthew Miller for sharing data from the National Firearms Survey. We also thank Marcella Alsan, Desmond Ang, Jurgen Brauer, David Johnson, Zane Kashner, Jack McDevitt, Matthew Miller, Phil Cook, and Jeffrey Smith for their helpful feedback on this paper, as well as seminar audiences at the NBER Conference on Firearm Markets, Boston University, Georgetown, BSE Summer Forum, CELS, Harvard University, the Advances in Industrial Organization Conference, and SITE Public Economics. This project was supported in part via Arnold Ventures Grant#23-10007. This paper was previously circulated as "The U.S. Consumer Firearms Industry: Welfare and Policy Implications".

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

‡Assistant Professor of Economics, Graduate School of Business, Stanford University. email: adam.rosenberg@stanford.edu

1 Introduction

Regulating an industry frequently requires buy-in from stakeholders, such as consumers and their advocacy organizations. In the U.S., gun regulation has stagnated due to political opposition from these groups (Lacombe 2021). The federal government levies an excise tax on firearms using rates that have not changed since 1954 (Congressional Research Service 2023). In the seven decades since these rates were last set, the U.S. firearm industry has changed substantially: manufacturers have consolidated into strategic, multi-brand conglomerates that sell millions of firearms each year, as consumers have accumulated a firearm stock that outnumbers the U.S. population (Brauer 2013, Azrael et al. 2017, Karp 2018). This large firearm stock, paired with lax regulations, has been consequential for public health and crime. As of 2020, firearm injuries are the leading cause of death among minors in the U.S. (Lee et al. 2022). Despite their welfare implications, federal regulators lack evidence on how to design consumer firearm taxes that are both politically feasible and economically efficient (Smart 2021).

This paper studies how market power and tax design determine the allocation of surplus in the U.S. consumer firearms industry. To guide our analysis, we derive theoretical results on efficient commodity taxation under a constraint representing political opposition to policies that allocate insufficient surplus to consumers. We then develop and estimate an equilibrium model of the firearms industry, from which we predict counterfactual prices, quantities, and allocations of surplus under alternate tax systems. We pair our predictions of new firearm purchases with a calibrated model of the total stock of firearms—and their effect on gun homicides—to quantify a key externality of the firearm industry. Our estimates suggest that the existing tax system is inefficient: firms set similar markups for handguns and long guns, but handguns cause significantly more homicides. To correct this misalignment, we simulate a policy that shifts taxes towards handguns and approximates our second-best tax system. This tax system would prevent 350 firearm homicides between 2016–2022; generate similar levels of industry profit, federal tax revenue, and consumer surplus to the status quo; and shift surplus towards consumers in areas with support for Republican politicians and the National Rifle Association ("NRA"). We also propose a simpler tax system that achieves similar gains with only two tax instruments.

Section 2 presents our theoretical results. We analyze a social planner that sets commodity taxes to maximize welfare subject to a constraint, that any tax change leaves consumers at least as well off as under the status quo. We interpret this constraint as a stylized representation of the political economy surrounding firearm regulation, where consumer interest groups influence policy outcomes (Lacombe 2021). Our analysis of the planner's problem

reveals necessary conditions, characterizing the constrained-optimal "second-best" firearm taxes. This tax system is different from prior results on optimal taxation in markets with negative externalities and imperfect competition (Buchanan 1969, O'Connell and Smith 2024, Conlon and Rao 2023), in that it requires lower taxes on products from which consumers derive more surplus. By the Ramsey (1927) rule, these high-surplus products are (roughly) those with less elastic demand. Thus, the constrained-optimal tax system both preserves the surplus enjoyed by inelastic, inframarginal consumers, while redirecting marginal consumers towards products that cause less external harm. In this system, tax rates depend on equilibrium relationships regarding consumers' substitution patterns, firms' markups, and transactions' external harms. These are the focus of our empirical analysis.

The key empirical challenge to our analysis is that few datasets match firearm purchase quantities to prices in the U.S. (Moshary et al. 2025), which we overcome by combining information from a variety of sources, described in Section 3. We access microdata on firearm purchases in Massachusetts between 2016-2022, recording all guns sold by firearm dealers, each gun's characteristics, and information about its buyer, as discussed in Armona and Rosenberg (2024). We proxy for aggregate firearm purchase quantities in the rest of the U.S., using data on federally-mandated firearm background checks, adjusted for state-year heterogeneity in background check policies (Brauer 2013, Smucker et al. 2022). Neither dataset contains prices, which we measure by linking each firearm model to its manufacturing dates and yearly MSRPs, recorded in an industry publication.[1]

With these data in hand, Section 4 begins our empirical analysis by estimating an average price elasticity of firearm demand, informing how a regulation that changes firearm prices may shift consumers' firearm purchase quantities. We specify this elasticity as the coefficient from a regression of log purchase quantity in Massachusetts onto log price, in a yearly panel of firearm models. In addition to conditioning on fixed effects, we address the potential endogeneity between firearm prices and quantities with a price instrument. Our instrument constructs cost shocks to firearm manufacturing, leveraging cross-manufacturer heterogeneity in exposure to the world commodity prices of metal inputs to firearm production (McDougal et al. 2023). We estimate a price elasticity of firearm demand of -2.1 for consumers in Massachusetts between 2016–2022, suggesting they are reasonably responsive to price changes. Because this elasticity is local to an average firearm model in Massachusetts, we require a richer specification to analyze how a price change would divert national demand across different firearm models, or away from firearm purchase altogether.

To further understand how consumer behavior and firm pricing interact with the design

[1] At firearm dealers, consumers face prices at, or very near to, MSRP (Moshary et al. 2025, Shapiro et al. 2026).

of firearm regulation, Section 5 develops and estimates an equilibrium model of demand and supply in the firearm industry. We specify demand for heterogeneous firearms as a random coefficients nested logit (Grigolon and Verboven 2014), where consumers' preferences depend on their demographics, a firearm's characteristics, and unobserved demand shocks. Our estimates reveal limited substitution between handguns and long guns, as well as systematic heterogeneity across consumer geography, race, gender, and voting patterns, in line with survey reports (Azrael et al. 2017, Parker et al. 2017). In aggregate, we estimate that the firearms industry generates $24 billion of annual surplus for U.S. consumers, around $90 per adult.

On the supply-side, our model assumes that each multi-product firm sets prices to maximize static profits from its production of firearms, under the federal excise tax on firearm sales. When a firm sets prices, it faces competition from both the new firearms produced at other firms and the used firearms on the second-hand market. Our estimates imply similar average markups for handguns and long guns, and reveal considerable pricing power. On average, 43% of a firearm's transaction price goes towards firm profits.

Section 6 analyzes firearm homicides, a relevant non-market externality of the consumer firearms industry.[2] We specify and calibrate a stylized model, linking the quantity demanded of new firearms to the evolution of the firearm stock and the expected flow of firearm homicides. Our calibrated model captures a key fact—handguns and long guns have roughly equal shares in the firearm stock, but handguns account for 90% of firearm homicides (Azrael et al. 2017, Moshary et al. 2025)—implying greater external harm from a marginal handgun purchase. Quantitatively, we find that the average purchase of a new handgun causes $4,881 of net present harm over its usable lifetime (1 lifetime homicide per 450 purchases), while the analogous long gun purchase would cause only $526 of net present harm (1 lifetime homicide per 4,200 purchases).[3] These quantities are robust across alternate specifications of our homicide model, including one that allows more dangerous individuals to select into handgun purchase. When measured by homicide fatalities, our model suggests that handguns cause most of the external harm from the firearms industry, and that handguns should be the target of corrective taxation.

Section 7 uses our estimated models to simulate counterfactual tax systems inspired by our theoretical results on politically-feasible optimal taxation. We simulate three tax changes: an approximation to the politically-feasible optimal tax from our theoretical analysis; a simpler two-tax policy in which the planner holds consumer surplus constant, while

[2]We account for neither incidents of self-harm, nor broader externalities (e.g., social well-being, mental health, and schooling) in measuring externalities caused by the level and composition of the firearm stock (Pienkny et al. 2024, Cabral et al. 2021).

[3]We construct these measures using a value of a statistical life of $9.5M (Peterson et al. 2021).

only "tilting" the relative tax rates on handguns and long guns; and a homogeneous 11 percentage point increase in the excise tax on all firearms, mirroring several state-level policy proposals. All three counterfactual tax changes improve social welfare in the U.S., with especially large gains over the medium-to-long run.[4]

Our counterfactual handgun-only taxes generate a useful targeting gain, shifting the tax burden towards socially harmful handguns and away from long guns. Although designed to improve efficiency, this targeting also has redistributive properties that may bolster its political feasibility. Because long guns are disproportionately valued by consumers in conservative neighborhoods, a handgun-only tax shifts surplus towards areas with more support for Republican politicians and the NRA, groups that have historically opposed firearm regulation (Bouton et al. 2021). In terms of implementation, we view the "tilted" handgun tax as a promising policy option, changing only two tax rates, but achieving 80% of the welfare gains from our approximation to the constrained-optimal tax schedule.

Our work relates to a growing literature on firearm markets (Koper and Roth 2002, Bice and Hemley 2002, Cook et al. 2007, Knight 2013, McDougal et al. 2023, Johnson and Robinson 2021, Hüther 2023, Bollman et al. 2025, Drango et al. 2025, Shapiro et al. 2026) and is especially close to two recent papers. Moshary et al. (2025) use stated-preference data from a survey experiment to develop and estimate a model of firearm demand in product space, accounting for heterogeneous preferences over eighty firearm models and their hypothetical prices. Rosenberg (2025) uses administrative data from California to estimate a model of consumer preferences for an undifferentiated handgun, linking these preferences to the public health costs of handgun ownership. Both papers focus on the demand side of the licit firearms market. In contrast, we estimate an equilibrium model of demand and supply, which allows us to study how supply-side conduct shapes the design of firearm taxes, a topic that has been understudied due to limited data and policy variation (Smart 2021).

We also contribute to a broader literature on the interaction of market power and externalities, originating with Buchanan (1969). In settings with homogeneous products, market power offsets a negative externality by restricting equilibrium quantity, as in oil production (Asker et al. 2024), electricity (Mansur 2007), transportation (Brueckner 2002), and tobacco (Bulow and Klemperer 1998). However, in settings with differentiated products, market structure interacts with heterogeneous externalities and substitution patterns, creating additional scope for corrective taxation (Durrmeyer and Samano 2018, Reynaert 2021, O'Connell and Smith 2024, Barahona et al. 2023, Conlon and Rao 2023). We contribute to

[4]We measure welfare as the sum of consumer surplus, manufacturer profits, federal tax revenue, and (prevented) firearm homicide fatalities. We value each firearm homicide fatality that does not occur according to a statistical value of a life.

this literature by studying corrective taxation if policy cannot lower consumer surplus below the status quo, a constraint we interpret as a stylized representation of political resistance to gun policy in the U.S. Under this constraint, we show that federal firearm taxes can be re-targeted to reduce homicides, while maintaining consumer surplus on aggregate and shifting surplus towards consumers in segments that we expect to be politically pivotal.

2 Designing Firearm Taxes

To guide our study of firearm taxes, we derive an optimal commodity tax for an industry with distortions from market power and externalities. Our optimal tax problem considers both economic efficiency and political feasibility, motivated by the heavy politicization of firearm policy in the U.S. (Bouton et al. 2021). We examine many facets of political feasibility in our empirical analysis, but to maintain tractability, we focus our theoretical analysis on a single stylized representation of consumers as a key political bloc. In particular, we analyze a welfare-maximizing social planner whose tax schedule is constrained to preserve sufficiently high surplus among firearm consumers, following the Second Amendment "right of the *people* to keep and bear Arms" (emphasis added) and evidence that interest groups of firearm owners (e.g. the NRA) are an organized and dominant political force in this industry (Boehmke et al. 2013, Laschever and Meyer 2021). We refer to the resulting optimal tax system as "second-best," since it maximizes welfare only under the pre-existing distortions in the industry and the further constraint on political feasibility (Lipsey and Lancaster 1956, Buchanan 1969).

The planner, working in partial equilibrium, seeks to maximize welfare generated by the firearm industry, in a static fashion. Welfare is the sum of consumer surplus CS, producer surplus (profits) Π, government tax revenue $\mathcal{G}$,[5] and a negative externality Φ, representing the harm caused by economic activity in the industry:

$$\mathcal{W} = CS + \Pi + \mathcal{G} - \Phi. \tag{1}$$

In our empirical analysis, Φ represents dollars of social harm from homicides caused by the firearms industry. This approach could apply to other industries where economic activity generates profits and externalities (e.g., pollution from fossil fuels, innovation from drug discovery).

As policy instruments, the planner controls the commodity tax system, setting specific

[5]Our main analysis considers government revenue and private surplus equally valuable in social welfare. Appendix A extends to the case where the marginal value of public funds is not necessarily equal to one.

excise taxes $\tau_{j,t} \in \mathbb{R}$ on each product j in each market t.[6] If a firm charges price $p_{j,t}$, it receives revenue $p_{j,t} - \tau_{j,t}$. With marginal cost $c_{j,t}$, the firm's post-tax profit margin is $\mu_{j,t} \equiv p_{j,t} - \tau_{j,t} - c_{j,t}$.

The planner's feasible taxes $\vec{\tau} = (\tau_{j,t})_{j,t}$ are constrained by the economic welfare of firearm consumers as a bloc $CS(\vec{\tau})$. As our stylized representation of political feasibility, we assume the planner can implement any tax system that induces consumer surplus no lower than some exogenous benchmark $CS(\vec{\tau}) \geq CS_0$. When the benchmark represents status quo consumer surplus before a policy change $CS_0 = CS(\vec{\tau}_0)$, as it does in our empirical analysis, this constraint is equivalent to the consumer welfare standard from antitrust economics.[7]

The problem of choosing taxes that maximize welfare $W(\vec{\tau})$ under our political feasibility constraint $CS(\vec{\tau}) \geq CS_0$ resembles the classic Ramsey (1927) problem, extended to externality-producing commodities (Sandmo 1975). Here, the planner ensures that consumers, rather than a government or firms, achieve surplus above a benchmark. The solution to the problem is given below.

Proposition 1. *Suppose the social planner chooses specific excise taxes $\tau_{j,t}$ to solve the following optimization problem:*

$$\begin{aligned} \max_{\tau_{j,t}} \quad & \mathcal{W}(\vec{\tau}) \\ s.t. \ & CS(\vec{\tau}) \geq CS_0, \end{aligned} \tag{2}$$

with welfare as defined in Equation (1). Then the constrained-optimal specific tax on product j in market t satisfies:

$$\tau^*_{j,t} = \phi_{j,t} - \mu_{j,t} - \lambda \cdot \mu^M_{j,t} \tag{3}$$

Where $\phi_{j,t}$ is the magnitude of the (negative) marginal externality of product j in market t, $\mu_{j,t}$ is the profit margin charged by the firm, $\mu^M_{j,t}$ is the margin that would be charged under static profit-maximizing monopoly ownership of all products, and λ is the shadow cost (in welfare units) of the consumer surplus constraint $\partial \mathcal{W}(\vec{\tau}^*)/\partial CS_0 \geq 0$. All terms on the right-hand side are implicit functions of the tax $\vec{\tau}$.

Proof. See Appendix A. □

[6] Appendix A considers *ad valorem* taxes $\vec{v}$. The *ad valorem* and specific tax systems have similar second-best solutions, with $\vec{\tau}^* = \vec{v}^* \odot \vec{p}(\vec{v}^*)$. The exposition for specific taxes is slightly shorter.

[7] This representation of political feasibility also has a microfoundation in political economy. Suppose (i) all firearm consumers are members of a political coalition with veto power over federal legislation, similar to the NRA; (ii) this coalition can make zero-sum transfers among its member; and (iii) each consumer's utility is quasi-linear between their income and private surplus from the firearm industry, while placing no value on tax revenue, as we assume in Section 5. Then, the political coalition supports a tax system $\vec{\tau}$ if and only if $CS(\vec{\tau}) \geq CS(\vec{\tau}_0)$.

Our second-best tax formula generalizes several standard cases. In the first-best, when there are no distortions in the firearms industry $\vec{\mu} = \vec{\phi} = 0$ and no political constraints $\lambda = 0$, there is no partial-equilibrium justification for intervention, and optimal taxes are zero $\tau_{j,t}^{F.B.} = 0$. If policy is unconstrained $\lambda = 0$ and the industry generates externalities $\vec{\phi} \neq 0$ under perfect competition $\vec{\mu} = 0$, then the optimal tax equals the marginal external harm $\tau_{j,t}^{Pigou} = \phi_{j,t}$ (Pigou 1924). Whereas, the unconstrained $\lambda = 0$ solution with both market power $\vec{\mu} \neq 0$ and externalities $\vec{\phi} \neq 0$ equates the gross of tax profit margin to the marginal external harm $\tau_{j,t}^{Buchanan} + \mu_{j,t} = \phi_{j,t}$. As in Buchanan (1969), the relationship between market power and external harm determines the potential benefits from regulation in an industry.[8]

These standard cases differ from the second-best solution in Equation (3) when the constraint on political feasibility binds $\lambda > 0$. "Second-best" taxes should be lower overall, limiting the loss of consumer surplus, in proportion to the shadow cost of the constraint λ. Moreover, for products j, j' with similar externalities $\phi_{j,t} = \phi_{j',t}$ and margins $\mu_{j,t} = \mu_{j',t}$, but different monopoly margins $\mu_{j,t}^{M} < \mu_{j',t}^{M}$, second-best taxes should be higher on the product with lower margins under monopoly $\tau_{j,t}^{*} > \tau_{j',t}^{*}$.

Second-best taxes induce additional cross-product targeting that follows an "inverse" Ramsey rule. The optimal tax is *higher* on products with lower margins under monopoly, and hence with *more elastic* residual demand. Concretely, if there is no substitution across products, a multi-product monopolist sets margins equal to the inverse semi-elasticity of demand $\mu_{j,t}^{M} = q_{j,t}/(\partial q_{j,t}/\partial p_{j,t})$. When products are substitutes, the monopoly margins adjust, lowering prices on products that have many close substitutes and price-sensitive consumers. Since monopoly margins enter negatively, Equation (3) sets higher taxes on products with more elastic demand, from which many purchasers derive little surplus and would be willing to substitute on the margin. In contrast, taxing products with inelastic demand, for which many purchasers are inframarginal, would erode consumer surplus without changing purchase quantities. In essence, the politically-constrained optimal tax formula suggests that policymakers target higher taxes towards socially costly products that provide less benefit to consumers.

Our formula clarifies that, in order to efficiently tax firearms, policymakers need to understand three key objects: (i) the marginal social harm from each firearm purchase, (ii) the price-cost margins charged by firms in equilibrium, and (iii) the counterfactual margins that would be charged under monopoly conduct, as a measure of consumer valuations. In Section 7, we use estimates of these objects to construct alternate firearm taxes, to simulate the

[8] Unconstrained $\lambda = 0$, the magnitude of the optimal tax system increases in $\sum_{j,t} |\phi_{j,t} - \mu_{j,t}|$, and market power corrects all externalities without government intervention in the knife-edge case where $\vec{\mu} = \vec{\phi}$.

effects of these and other counterfactual policies, and to evaluate measures of their political feasibility, beyond the stylized notion in Equation (2). Sections 3–6 describe our approach to estimating these policy-relevant objects.

3 Setting and Data

This section describes key regulations on the U.S. firearms industry and the data elements used in this paper.

3.1 The Firearm Industry and Regulation

Firearms are popular consumer products in the U.S. Consumers purchase millions of firearms each year, and 40% of U.S. households owning at least one firearm (Berrigan et al. 2022). Consumers typically acquire firearms from retailers (e.g., Bass Pro Shops), who themselves are typically served by intermediate dealers (e.g., Davidson's Incorporated). Intermediate dealers buy from various firms upstream (e.g., Remington Outdoor Company), which may own multiple firearm manufacturers (e.g., Remington and Marlin). In this paper, we focus on the decisions of consumers, firms, and manufacturers abstracting from other components of the supply chain.

Firearms—having been linked to crime and violence (Duggan 2001, Cook and Ludwig 2006)—are also the subject of many proposed and implemented regulations. Excise taxes on newly-produced firearms were first set at 10% as part of the Revenue Act of 1918, statutorily paid by manufacturers at point of sale (Congressional Research Service 2023). These taxes were raised to 11% through the Revenue Acts of 1940 and 1941. The Excise Tax Reduction Act of 1954 set the tax on handguns to 10% and on long guns to 11%, and these rates have not been adjusted since. Many states require that consumers pay a one-time fee prior to their first in-state firearm purchase (e.g., New Jersey) for a license to carry firearms. Beyond taxes, there are restrictions on the types of firearms that may transact in the U.S. The 1934 National Firearms Act placed federal restrictions on the ability of consumers to own fully-automatic firearms (i.e., machine guns), later extended to other destructive devices (e.g., grenades). From 1994–2004, the Federal Assault Weapons Ban prevented the purchase of newly produced "assault weapons" (e.g., semi-automatic long guns like the AR-15). Although no longer federal law, several states, including Massachusetts, maintain such assault weapon bans.

Federal law passed in 1993—the Brady Act—also mandated background checks, beginning in 1998, on certain firearm purchases. These background checks are implemented by the

FBI under the National Instant Criminal Background Check System ("NICS"). States are given some leeway in the particular design of their background check system, with several states, including Massachusetts, mandating background checks before every firearm transfer.

Firearm markets are segmented at the state level. The Gun Control Act of 1968 banned transaction of firearms across state lines except by individuals with a Federal Firearms License ("FFL"), the legal vendors of firearms. Federal law sets minimum standards for registration as an FFL, with states able to set more stringent requirements. In addition to the regulations above, states and localities also differ in their policies on how firearms may be purchased, where firearms may be carried, and the ways firearms may be used.

3.2 Market size and consumer demographics

We define a market for legal consumer firearms by state-year. Within a state, we measure the size and demographic characteristics of each zip code's adult population using the 5-year estimates from the (2015-2019) American Community Survey. We also measure each zip code's conservative vote share in the 2016 Presidential Election (Martin and Yurukoglu 2017).[9] As a firearm-specific measures of political engagement, we collect data from the Federal Elections Commission on donations to the NRA in each zip code, and data from a digital archive on the NRA's grading of U.S. House members in 2018, based on their attitudes and voting records on gun policy.[10]

Some adults within each zip code may have no interest in owning a firearm, a fact incorporated into our definition of consumer market size (Stantcheva 2025). We measure those potentially interested in firearm ownership using a Pew survey of 3,900 U.S. adults from 2017, linking respondents' stated personal firearm ownership and interest in becoming a firearm owner with their demographic and political characteristics (Parker et al. 2017). Our procedure predicts the market size of potential gun purchasers ("consumers") in each zip code by regressing a binary indicator for each respondent's (interest in) firearm ownership onto their demographic characteristics.[11] The procedure is described in full in Appendix B.2. We use the coefficients from this regression to linearly transform each zip code's overall adult population into the subset of its adult population open to firearm ownership. We also use the demographic data on zip codes to estimate demographic heterogeneity in the demand

[9] We allocate the precinct-level vote shares from Voting and Election Science Team (2018) onto zip codes based on geographic overlap of the regions, weighted by total number of votes.

[10] NRA donation data is downloaded from https://www.fec.gov/. NRA grade data was collected from https://www.everytown.org/nra-grades-archive/

[11] This indicator equals one for respondents who either state that they "own a gun" or state that they "could see themselves owning a gun." It equals zero for those who state they could "never see themselves owning a gun."

for firearms.

Our procedure measures the consumer market size for legal firearms as about two-thirds of the adult population. Figure B.1 shows that men, U.S. citizens, and Republicans are over-represented among consumers in the firearm market.

3.3 Microdata on firearm purchases within Massachusetts

Within Massachusetts, we use administrative microdata on firearm transactions involving FFLs as recorded by the state's Firearms Record Bureau (FRB), representing the universe of firearm purchases from dealers.[12] Each time a consumer purchases a firearm in Massachusetts, state law requires the seller to perform an electronic verification of the consumer's firearm purchase license. Our data is the publicly available list of all such verifications by FFLs in Massachusetts.[13] With our focus on behavior in the new firearm market, we do not use the FRB's information on the less-frequent sales of firearms by private individuals, as sales by private individuals are definitionally sales of used firearms.[14]

Each transaction in these data is tagged with additional information. The FFL manually inputs information about the firearm itself, including its manufacturer (e.g., Colt), model (e.g., Python), firearm class (handgun or long gun), and physical characteristics (caliber, barrel length, and a flag for high capacity magazine). The FRB provides information about the consumer—including their gender and residential zip code—by merging to their state-issued firearm license. The FRB verification system itself records the transaction date.

We convert recorded calibers to a standard unit of inches across all firearms. We perform a similar standardization for barrel length. We manually standardize the manufacturer field for any manufacturer with at least 30 purchases in these data. Although the FRB began recording these data in 2006, they began tracking whether firearms are high capacity in 2016, so we focus on the period 2016–2022 to include this variable in our analysis.

Table 1 shows how demographics of the Massachusetts adult population compare to the characteristics of consumers in the legal firearm market and actual gun purchasers. Relative to the average adult, gun purchasers are more likely to be male, and live in less racially diverse, more rural, and more politically conservative neighborhoods. These differences are less pronounced among handgun purchasers and more pronounced among long gun purchasers.

Appendix Figure B.2 shows that the handgun share of firearm transactions has been

[12] Other papers using these data include Braga and Hureau (2015), Iwama and McDevitt (2021), Balakrishna and Wilbur (2022), and Johnson et al. (2023).

[13] Available at https://www.mass.gov/info-details/data-about-firearms-licensing-and-transactions

[14] Massachusetts state law allows private individuals to sell four or fewer firearms each calendar year (MA G. L. c. 40 §128A)

roughly constant at 65% from 2016–2022, while the share of new firearms transacted declined from 80% to 60% over this same period. The timing of the shift towards used guns in the market is consistent with anecdotal evidence from the COVID-19 pandemic. Consumers turned to the secondary market as supply chain issues limited the ability of manufacturers to keep up with the surge in demand for firearms (Miller et al. 2022).[15]

A limitation of the FRB data is that they do not record the prices at which firearms transact. We are unaware of any dataset recording both firearm prices and transaction quantities, a gap noted by Moshary et al. (2025).

3.4 The Blue Book of Gun Values and firearm prices

For information on firearm prices, we use historical data from the Blue Book of Gun Values ("BBGV"), an industry pricing guide. For each manufacturer-model, BBGV records the years in which it was actively produced and each year's manufacturer suggested retail price (MSRP). As prices quoted by firearm retailers across the U.S. closely follow MSRP (Moshary et al. 2025), we assume that new firearms are priced nationally and that the price facing consumers is the MSRP.

A challenge we face is that the names of firearm manufacturers and models are manually entered in the FRB dataset, and so are neither standardized across records, nor to the conventions in BBGV. We link recorded models in the FRB to their standardized BBGV models via a string matching algorithm described in detail in Section B.4. We are able to match 90% of purchases in the FRB to a unique manufacturer-model in BBGV. Throughout our analysis, we assume that if a firearm was purchased from an FFL during a year in which it was under active production, then it was the purchase of a new (not used) firearm. This aligns our definition of used firearms with empirical estimates: the 2019 National Firearms Survey records that used guns accounted for 26.8% of all firearm purchases from 2017–2019, and our own metric based on active production implies 27.1% of all firearms purchased in Massachusetts during this period were used (Berrigan et al. 2022).

Table 2 displays summary statistics for the firearm models we match between BBGV and FRB. Handguns typically have a shorter barrel length and are lower in caliber than long guns (rifles and shotguns), though Figure B.3 shows there is considerable heterogeneity in barrel length and caliber across firearms, even within the same class. We also show transaction-weighted averages, which reveal that Massachusetts' consumers tend to purchase guns that are cheaper (relative to those available), and more likely to be high-capacity, and more likely to be under active production. The average price of a new firearm purchase in Massachusetts

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

is around $820. Handguns are typically cheaper than long guns, though there is significant dispersion of prices within class. Figure B.4 displays the evolution of prices within a firearm model from 2016–2022. Prices remained roughly constant from 2016–2020, before increasing by 8% from 2020-2022 during the COVID-19 pandemic.

3.5 Auxiliary Datasets

We use a number of other datasets as input to our analysis.

Aggregate firearm purchase quantities To measure firearm demand in U.S. states besides Massachusetts, we estimate the aggregate quantity of firearm purchases in each state-year. We form these estimates by developing and applying an adjustment to publicly available data on NICS background checks. Background checks are a common proxy for firearms purchases (Brauer 2013), but may suffer from reporting biases due to state-year differences in background check requirements (Smucker et al. 2022). For example, in some states, firearm permit holders are exempt from background checks when purchasing additional firearms, which would bias our estimates of total purchases downwards. We account for these differences in our adjustment procedure, which is described in detail in Section B.3.

Firearms stock Because firearms are durable goods, our analysis of firearm homicides requires a measure of total guns in circulation. We measure the overall U.S. firearm stock, using published estimates of the count of firearms in circulation in the U.S. in 1994 and 2015 (Cook and Ludwig 1996, Azrael et al. 2017). In the intervening years, we measure historical flows of firearms into circulation using administrative estimates of firearms produced in the U.S. for consumers—newly manufactured, non-military domestic firearms, less net exports—from the Bureau of Alcohol, Tobacco, Firearms and Explosives ("ATF").[16]

We also bring together information on the distribution of the firearm stock across U.S. households. For this information, we gather data on the count of households per state from the 2015 ACS, the share of U.S. adults living with a firearm in their household in 2015 by state from Schell et al. (2020), the average number of firearm owners living with each firearm-owning adult in 2015 from Azrael et al. (2017), and the average number of firearms owned by each firearm owner in 2015 by census region and firearm class from Azrael et al. (2017).

Firearm homicides Our analysis of firearm homicides uses standard sources of publicly available state-year data from 2016–2022: counts of firearm homicides from the Current

[16] ATF data is downloaded from https://www.atf.gov/resource-center/data-statistics.

Final Multiple Cause of Death Data from the National Center for Health Statistics, and the share of firearm homicides by firearm class (or unknown) from the National Violent Death Reporting System.[17]

Ownership of firearm manufacturers We manually reconstruct a panel of entities owning each of the top 99 firearm manufacturers by sales volume in Massachusetts. Our reconstruction uses information on the historical merger and acquisition activities of these manufacturers, drawing primarily on histories from BBGV and the manufacturers' own websites. This procedure captures the fact that, although the FRB and BBGV datasets record firearm manufacturers, multiple manufacturers may be wholly owned by a single parent company.[18] This ownership panel covers 94% of FRB transactions during our sample period, and 91% of total production in the ATF data.[19]

Figure B.8 shows the share of U.S. firearm production accounted for by the top-four parent companies each year. The four largest firms account for about half of total firearm production each year, and around two-thirds of production within each firearm class. Appendix B.1 documents additional supplementary data sources used in the analysis.

4 Pricing in Firearm Markets

In this section, we measure a cost shock to firearm manufacturing. We use this shock as a price instrument, to study the price elasticity of consumer demand for firearm models.

4.1 Constructing cost shocks

Our analysis uses shocks to the input costs of firearm production as instruments for price, following suggestions from the literature on demand estimation (Berry and Haile 2021), and building on past research into the firearms industry (McDougal et al. 2023). Shocks to input costs are likely to be an effective price instrument in the firearms industry, where much of the costs of manufacturing are tied up in the purchase of materials—such as steel and aluminum—that transact on the world commodity market.[20]

[17] Data for both are downloaded from https://wisqars.cdc.gov/.

[18] For example, the Italian firearm manufacturer Beretta wholly owns the other Italian firearm manufacturers Benelli and Franchi.

[19] For manufacturers below the top 99, we assume that they are independent, self-owned entities.

[20] Intermediate inputs account for 65% of the costs of firearm manufacturing, with 35% of intermediate expenditures on primary metals (e.g. steel) and fabricated metal products (e.g. iron molds), per the U.S. Bureau of Economic Analysis (2017, excluding intermediate inputs from the "Ammunition, arms, ordnance, and accessories manufacturing" category).

As we do not observe the exact mix of inputs used to produce each firearm model, our strategy isolates cost shocks by flexibly approximating the production and pricing practices of different manufacturers (Villas-Boas 2007). Specifically, we assume that the following price regression holds:

$$\log(p_{j,t}) = \underbrace{\tilde{\gamma}_j}_{\text{Gun FE}} + \underbrace{\tilde{\eta}_{c(j),t}}_{\text{Class-Year FE}} + \sum_{k \in \mathcal{K}} \underbrace{\left(\beta_{m(j),k} \log(\tilde{p}_{k,t})\right)}_{\text{Cost shock}} + \epsilon_{j,t}, \tag{4}$$

in panel data on firearm j, of firearm class $c(j) \in \{\text{handgun}, \text{long gun}\}$, produced by manufacturer $m(j)$ during year t. The summation is over different inputs to firearm manufacturing k, each with world commodity price $\tilde{p}_{k,t}$.[21] Our preferred specification of inputs $\mathcal{K}$ is a subset of the 6-digit commodities in the PPI that begin with 10 (ferrous and nonferrous metals), yielding $|\mathcal{K}| = 42$ commodity price time series—from cold rolled steel sheets to aluminum scrap—that are potentially relevant for firearm manufacturing.[22] Figure C.1 shows considerable annual variation in the world prices of metal commodities k that we use for estimation. The largest shock to these input prices coincides with the COVID-19 pandemic over 2020–2022, which both caused global supply chain issues and drove U.S. demand for firearm purchase to historic highs (Sokol et al. 2021).

The parameter $\beta_{m,k}$ in Equation (4) governs the pass through of the world commodity price $\tilde{p}_{k,t}$ into the average firearm price $p_{j,t}$ for each manufacturer-commodity pair m, k. Since we include fixed effects for firearm model $\tilde{\gamma}_j$ and firearm class-year $\tilde{\eta}_{c(j),t}$, $\beta_{m,k}$ measures the *heterogeneous responsiveness* of manufacturers' prices to common input cost shocks, relative to our sample of manufacturers, residualizing price variation attributable to aggregate shocks such as COVID.[23]

We estimate $\beta_{m,k}$ using data on all firearm models under active production that have at least 100 purchases in Massachusetts from 2016–2022.[24] This gives us 422 unique firearm

[21] We measure commodity prices $\tilde{p}_{k,t}$ by aggregating the monthly Producer Price Index to annual frequency, with a base year of 2016. See https://download.bls.gov/pub/time.series/wp/wp.txt.

[22] To construct the set of inputs $\mathcal{K}$, we start from all 52 commodity price time series of ferrous and nonferrous metals, and hand remove 10 commodities that play a minimal role in firearms manufacturing: 3 plumbing commodities (ductile iron pressure pipes, cast iron pressure pipes, and copper pipes), 5 wire/cable commodities (electronic wires, copper wires, power wires, fiber optic cables, and other wire/cables), used beverage can scrap, and secondary precious metals (e.g., gold and silver). Secondary precious metals are sometimes used in firearm manufacturing, but their cost share is trivial.

[23] For example, if a manufacturer primarily uses cold rolled steel sheets in firearm production, they will be more exposed to the jump in the price of cold rolled steel sheets in 2021, relative to a manufacturer that uses other metals during production, and so would have a positive coefficient $\beta_{m,\text{cold rolled steel sheets}}$. The primary threat to identification would be if demand shocks, such as COVID-19, were correlated with the reliance of certain manufacturers on particular inputs. This could occur if the surge in gun demand during COVID-19 was driven by a surge in demand for guns specifically made from cold rolled steel sheets.

[24] This restriction is applied so that our log specification does not exclude firearm-years because sales drop

models produced by 66 manufacturers during our sample, accounting for 93% of new firearms purchases in Massachusetts. With 42 potential commodity inputs, this implies $\dim(\vec{\beta}) = 66 \times 42 = 2,772$ potential parameters to estimate, meaning our 2,464 model-year observations cannot fully identify the parameter $\vec{\beta}$. Even if we remove potential inputs to firearm production from $\mathcal{K}$, and achieve point identification, there may be remaining issues with weak instruments and small sample variability that pollute our estimates of pass through $\beta_{m,k}$ and limit their informativeness.

To efficiently use our data, we estimate the pass-through parameters $\beta_{m,k}$ with the dimension-reduction procedure of Belloni et al. (2012). This procedure first estimates Equation (4) via LASSO—using a data-driven penalty based on econometric theory and the assumption of independent but heteroskedastic errors $\epsilon_{j,t}$—to select a subset of parameters $\beta_{m,k}$ for inclusion in the model. Following model selection of $\beta_{m,k}$, we re-estimate Equation (4) via OLS, zeroing out the un-selected pass-through parameters. Heuristically, this procedure finds and estimates the subset of manufacturer-by-commodity pass-through parameters $\beta_{m,k}$ that can be most feasibly differentiated from one another in data. Although we use it to study metal commodities in firearm production, combining the high-dimensional model of cost shocks from Villas-Boas (2007) with the dimension-reducing estimator of Belloni et al. (2012) is agnostic to the specifics of firm production, and this procedure could be applied to other industries.

Figure C.2 displays our estimates of the selected pass-through parameters $\hat{\beta}_{m,k}$. From the potential set of 2,772 such parameters, 7 are selected, suggesting that many of the manufacturer-specific coefficients are uninformative about firearm pricing. The 7 selected pass-through parameters are associated with some of the largest firearm manufacturers, covering 40% of the firearm models in our sample. This ensures that the cost shocks we study impact a wide range of firearms across the industry.

Appendix C provides a battery of robustness checks on the LASSO selection procedure we use to construct cost shocks. Our results are not sensitive to the level of aggregation across metals inputs $\mathcal{K}$ (e.g., 2-digit or 6-digit commodities), nor to the penalization scheme we use in reducing the dimension of commodity-manufacturer passthrough parameters $\beta_{m,k}$. We also show that our estimated cost shocks $\beta_{m(j),k} \log(\tilde{p}_{k,t})$ do not anticipate price changes the following year $t+1$; are equally predictive of prices for in- and out-of-sample firearms; and cover manufacturer-commodity pairs in which the manufacturer m has a substantive business connection, and disproportionately many firearm patents based on metal k. These facts help establish the relevance of our constructed cost shocks as an instrument for firearm

to zero in a particular year. In our equilibrium model, we similarly restrict the choice set to be weapons above a threshold number of purchases.

prices.

4.2 Price elasticity of firearm demand

We now move to measuring a price elasticity of firearm demand among consumers in Massachusetts. Specifically, we estimate the equation

$$\log(q_{j,t}) = \gamma_j + \eta_{c(j),t} + \alpha \log(p_{j,t}) + \xi_{j,t}, \tag{5}$$

using the same firearm-year panel data as for our estimates of Equation (4). The dependent variable is the log purchase quantity of firearm j by consumers in Massachusetts during year t.[25] Analogous to Equation (4), we model firearm purchase quantity $\log(q_{j,t})$ as a function of the log of the firearm's own price $\log(p_{j,t})$, fixed effects for firearm model γ_j and weapon-class year $\eta_{c(j),t}$, and an unobserved demand shock $\xi_{j,t}$.

The parameter of interest is α, which measures an own-price elasticity of firearm demand. This elasticity captures changes in consumer behavior at the level of the firearm model j, combining adjustments in demand along both the extensive margin (buying no gun) and intensive margin (buying a different gun).

A challenge to estimating the elasticity of demand α is the endogeneity of price and quantity. In Equation (5), this behavior would manifest as a positive correlation between prices $p_{j,t}$ and demand shocks $\xi_{j,t}$, even conditional on the fixed effects, generating an upwards bias in the OLS estimator. To avoid this bias, we estimate the elasticity of demand $\hat{\alpha}$ by using our first-stage cost shocks $\sum_k \hat{\beta}_{m,k} \log(\tilde{p}_{k,t})$ to construct an instrument for firearm prices $\log(p_{j,t})$, applying the procedure of Belloni et al. (2012) for consistent estimation and asymptotic inference in the presence of first-stage model selection. This estimator may generate a different upwards bias due to the joint structure of the instrument: If a cost shock and price increase to firearm j would cause a consumer to substitute *ceteris paribus*, Equation (4) implies the same shock will cause a price increase for all other firearms j' in manufacturer $m(j)$'s portfolio, making the consumer's other options j' less attractive. This may attenuate the substitution patterns created by our instrument and generate a positive bias in our estimated own-price elasticity, though Table 3 shows this bias is less severe than from estimation via OLS.

Table 3 shows our preferred estimate of the own-price elasticity of firearm demand $\hat{\alpha} = -2.1$. The uncertainty measures from Belloni et al. (2012) also support the relevance of

[25] The log transformation drops 21 firearm model-years (0.85% of potential observations) from our sample. It is not defined for model-years with 0 purchases, even if the model is under active production and has positive sales in other years. Results are similar if we use the inverse hyperbolic sine transformation to preserve zeros.

our excluded instruments (via the sup-score test-statistic).[26] The standard error rejects that firearm demand is perfectly inelastic to own-price, and suggests demand is more than unit elastic.[27] We estimate more elasticity in firearm demand than the survey-based results of Moshary et al. (2025), though both approaches imply that consumers are price conscious in their choice of firearm model.

Our estimate of the own-price elasticity of firearm demand is local to the average firearm model spanned by our instruments, limiting its interpretation. Nor does this elasticity speak to which, if any, substitute firearm a consumer might purchase because of a price change; to the difference in substitution patterns across firearm models; or to price elasticities for consumers outside Massachusetts. The following section addresses these limitations by imposing additional structure via a model of supply and demand in the U.S. firearms industry.

5 An Equilibrium Model of the Legal Firearms Market

This section specifies and estimates a model of supply and demand in the U.S. firearms market. The model allows for rich substitution patterns, capturing heterogeneity in own- and cross-price elasticities between firearms and the no-purchase outside option, as well as heterogeneity along both observed and unobserved consumer characteristics, as in the random coefficients nested logit (RCNL) model of Grigolon and Verboven (2014). The model also allows us to measure the costs of firearm production and understand manufacturers' pricing behavior.

5.1 Demand

We model market demand for firearm model j, with price $p_{j,t}$, observable characteristics $\vec{X}_j$, and firearm class $c(j) \in \{h, l\}$ (handguns and long guns, respectively).[28] There are $\mathcal{J}_t^N$ firearm models under active production ("new") available each year t, constant across consumers.[29]

Demand during year t in state s arises from the discrete choice decisions of consumers

[26] Our sup-score test statistic is 12.55, well above the critical value of 4.56 needed for 1% significance when testing the relevance of multiple instruments.

[27] The Z-test against perfectly inelastic demand is $\Phi(2.1/0.8) \approx 0.99$, and against unit elastic demand is $\Phi((2.1-1)/0.8) \approx 0.9$, for $\Phi(\cdot)$ the standard normal CDF.

[28] Observable characteristics X_j are a vector of caliber, barrel length, an indicator for high-capacity, indicators for long guns and shotguns, and a constant.

[29] The set $\mathcal{J}_t^N$ contains all firearm models under active production during year t that have at least 50 observed transactions in Massachusetts and a price below $2,500. These represent about 60% of all transactions, and about 85% of new firearms transactions. We group all other firearms into a composite outside good, as described below.

i with observable demographics $\vec{D}_i$.[30] There are M_s consumers in state s, with both mass and demographic distribution constant across years t.

Indirect utility to consumer i from purchasing new firearm $j \in \mathcal{J}_t^N$ is

$$\begin{aligned} u_{i,j,s,t} &= \overbrace{-\alpha_i \cdot p_{j,t}}^{\text{price}} + \overbrace{\vec{\beta}_i \cdot \vec{X}_j}^{\text{characteristics}} + \overbrace{\delta_{j,s,t}}^{\text{fixed effects}} + \overbrace{\epsilon_{i,j,s,t}}^{i.i.d.\text{ shock}} \\ u_{i,0,s,t} &= \epsilon_{i,0,s,t}, \end{aligned} \tag{6}$$

with $u_{i,0,s,t}$ the indirect utility of choosing to not purchase a firearm.

The first term is consumer i's disutility from the price of a firearm $p_{j,t}$. We assume that the price coefficient α_i is i.i.d., following the log-normal distribution

$$\log(\alpha_i) \sim N(\alpha + \vec{D}_i \cdot \vec{\Pi}_\alpha, \sigma_\alpha).$$

The mean is governed by scalar α and vector of demographic shifters Π_α, and the variance is σ_α^2. The log-normal specification ensures that all consumers dislike price, or $\partial u_{i,j,s,t}/\partial p_{j,t} < 0 \quad \forall i, j, s, t$.

The second term governs consumer i's preferences over the observable characteristics of firearms $\vec{X}_j$. We assume that these preferences β_i are i.i.d., following the multivariate normal distribution

$$\vec{\beta}_i \sim MVN(\vec{\beta} + \vec{D}_i \cdot \Pi_\beta, \Sigma_\beta).$$

The mean is governed by the vector β and the matrix of characteristic-by-demographic shifters Π_β. Throughout, we assume that conditional on observable demographics D_i, the price coefficient α_i and each element of the preferences for observable characteristics β_i are mutually independent, a consequence of which is that the variance-covariance matrix Σ_β is diagonal.[31]

We specify the mean utilities $\delta_{j,s,t}$ via fixed effects across several hierarchies:

$$\delta_{j,s,t} = \delta_j + \tau_{s,t} + \eta_{c(j),t} + \xi_{j,t}. \tag{7}$$

The first term δ_j adjusts for time invariant characteristics of firearm j that may not be observed in data (e.g., the quality of its firing mechanism). The second term $\tau_{s,t}$ captures

[30] Observable demographics $\vec{D}_i$ are a vector of an indicator for female and—measured for i's zip code—the percent white, percent conservative, log median household income, and log population per square mile.

[31] We allow random coefficients on barrel length, caliber, and indicators for shotgun and high capacity weapon.

vertical differences in the taste for firearms across markets (state-years). The third term $\eta_{c(j),t}$ adjusts for changes over time in the preference for firearms, which may differ by class-year. The final term $\xi_{j,t}$ reflects other time-varying determinants of aggregate demand for firearm j (e.g., a national advertising campaign). Notably, the only component of the fixed effect that varies across states is $\tau_{s,t}$, which is also constant across firearms. We normalize this state-year shock $\tau_{s,t}$ to zero in Massachusetts each year.

The remaining term $\epsilon_{i,j,s,t}$ is an i.i.d. shock across consumers, which we assume follows a three-level nested logit (Train 2009). Within a consumer-state-year ist, our specification of $\epsilon_{i,j,s,t}$ allows for a flexible correlation structure between handgun and long gun models, as well as between the purchase of any firearm and the no-purchase outside option. Specifically, we assume:

$$\epsilon_{i,j,s,t} = \zeta_{i,0,s,t} + \rho_0 \zeta_{i,c(j),s,t} + \rho_1 \tilde{\epsilon}_{i,j,s,t},$$

where $\zeta_{i,0,s,t}$ is a preference shock to the outside option, governing substitution between buying and not buying a firearm; $\zeta_{i,c(j),s,t}$ is a preference shock at the weapon-class c level controlling substitution between handguns and long guns; and $\tilde{\epsilon}_{i,j,s,t}$ a product-specific idiosyncratic preference shock distributed extreme-value type 1.[32] The parameter $\rho_0 \in [0,1]$ controls the correlation of preferences across all firearms, relative to the outside option, while parameter $\rho_1 \in [0, \rho_0]$ controls the correlation of preferences for firearms within a class. The shocks generate a standard i.i.d. multinomial logit distribution when $\rho_0 = \rho_1 = 1$.

As firearms are durable goods, consumers may also choose to purchase a firearm not under active production ("used") on the secondary market, which we model as a firearm class-specific composite good ω_c. Each consumer's choice set in year t is thus $\mathcal{J}_t = \mathcal{J}_t^N \cup \{\omega_c\}_c \cup \{0\}$. Indirect utility from a used firearm of class c is

$$u_{i,\omega_c,s,t} = \beta_{i,\omega} + \delta_{\omega_c,s,t} + \epsilon_{i,\omega_c,s,t}$$
$$\beta_{i,\omega} \sim N(\vec{D}_i \cdot \vec{\Pi}_\omega, \sigma_\omega),$$

analogous to new firearms in $\mathcal{J}_t^N$. Consumer i's taste for used firearms is i.i.d. normally distributed with mean governed by the parameter vector Π_ω; variance equal to σ_ω^2; and mutually independent from preferences for new firearms (α_i, β_i), conditional on their demographics $\vec{D}_i$. We extend the fixed effects $\delta_{\omega_c,s,t}$ and nested logit structure $\epsilon_{i,\omega_c,s,t}$ for used firearms using the same hierarchy as for new firearms.[33] For parsimony, the used-good composite ω_c contains

[32] The shocks $\zeta_{i,0,s,t}$ and $\zeta_{i,c(j),s,t}$ follow the unique distribution conjugate to the type 1 extreme value, such that $\epsilon_{i,j,s,t}$ is itself type 1 extreme value (Cardell 1997).

[33] Specifically, fixed effects are $\delta_{\omega_c,s,t} = \delta_{\omega_c} + \tau_{s,t} + \eta_{c(\omega_c),t} + \xi_{\omega_c,t}$. For the nesting structure, we place used

all firearms sold in Massachusetts in year t that are not under active production at any point (82% of transactions in ω_c), have fewer than 50 total purchases that year (16%), or are ever priced at \$2,500 or higher (1%), thus representing all other firearms purchased, and ensuring the model predicts observed aggregate demand for firearms.

5.2 Supply

We model the pricing decisions of firms f that wholly own one or more firearm manufacturers m.[34] Firms compete in static Nash-Bertrand fashion each year, simultaneously setting nationwide prices (MSRP) on the firearm models under active production by all their wholly-owned manufacturers $\mathcal{J}_{f,t}$. For each unit of firearm j sold, its owner $f(j)$ must pay the year-specific, constant marginal cost of production $c_{j,t}$ and the federal excise tax with an *ad valorem* rate v_j.

We assume that each firm sets prices to maximize its profits each year, treating as exogenous its set of firearms $\mathcal{J}_{f,t}$, the firearms of its competitors $\mathcal{J}_{f',t}$, and the evolution of these sets over time. We also assume that used firearms in ω_c are perfectly competitively supplied, such that firm pricing of new guns cannot affect the price of used firearms.

These assumptions imply the following problem of setting prices to maximize profits:

$$\max_{\vec{p}_{f,t}} \sum_{j \in \mathcal{J}_{f,t}} q_{j,t} \times (p_{j,t}(1 - v_j) - c_{j,t}), \tag{8}$$

where $\mathcal{J}_{f,t}$ is the set of new firearms produced by firm f during year t, $\vec{p}_{f,t}$ is its vector of prices, and $q_{j,t}$ is the yearly quantity of firearm j sold nationwide. This problem has $|\mathcal{J}|_{f,t}$ first order conditions for profit maximization each firm-year, which, if unique, exactly govern pricing behavior.

5.3 Estimation and Identification

This section provides an overview of the steps to estimate the parameters in our model and the data that identifies each parameter. Additional details and computational specifics are provided in Appendix D.

First, we jointly estimate the mean utilities of firearms in Massachusetts each year, $\delta_{j,MA,t} = \delta_j + \eta_{c(j),t} + \xi_{j,t}$, and the parameters governing preference heterogeneity $\Theta =$

firearms ω_c in their respective bottom-level nests, alongside new firearms of the same class.

[34] A firm f sets prices on all products it owns $f(j) = f(j') = f$, even if produced by different manufacturers $m(j) \neq m(j')$. Our analysis of firearm manufacturing in Section 4 allows separate production costs across manufacturers owned by the same firm.

$(\rho_0, \rho_1, \alpha, \Pi_\alpha, \sigma_\alpha, \Pi_\beta, \Sigma_\beta, \Pi_\omega, \sigma_\omega)$, using data on firearm purchases in Massachusetts and a constrained MLE procedure (Goolsbee and Petrin 2004, Train 2009). Through this procedure, the heterogeneous tastes for characteristics $(\Pi_\beta, \Sigma_\beta)$ and prices $(\Pi_\alpha, \sigma_\alpha)$ are identified by variation in purchasing patterns across demographic cells of Massachusetts' consumers.

To identify the mean utilities in Massachusetts $\delta_{j,MA,t}$, the mean taste for prices α, and the nesting parameters (ρ_0, ρ_1), we include the following constraints on the standard MLE objective:

$$0 = Pr_{s,t}(j) - \hat{s}_{j,s,t} \quad \forall t, j \in \mathcal{J}_t, \quad s = \text{MA} \tag{9}$$

$$0 = E[\xi_{j,t} \cdot \vec{Z}_{j,t}] \tag{10}$$

$$0 = E[\eta_{c,t} \cdot \tilde{J}_{c,t}] \quad \forall c \in \{h, l\}. \tag{11}$$

Equation (9) ensures that the model-implied, market-level probabilities of choosing each firearm $Pr_{s,t}(j)$ are equal to those observed in the microdata from Massachusetts $\hat{s}_{j,s,t}$, identifying the mean utilities $\delta_{j,MA,t}$ associated with each product. We then regress these mean utilities $\delta_{j,MA,t}$ on the set of fixed effects (except $\tau_{s,t}$, which is normalized to zero in Massachusetts) outlined in Equation (7) to recover estimates of δ_j, $\eta_{c,t}$, and $\xi_{j,t}$. We further project estimates of the product fixed effects δ_j onto firearm characteristics $\vec{X}_j$ to recover baseline tastes for characteristics $\vec{\beta}$ (Nevo 2000).

Equation (10) ensures that the unobserved product-year demand shocks are uncorrelated with a vector of two cost-shock instruments $\vec{Z}_{j,t}$. The first instrument is the cost shock from commodity metals described in Section 4.1. This instrument isolates exogenous variation in firearm prices, which serves to identify the mean of the log price coefficient α.

The second cost-shock instrument in $\vec{Z}_{j,t}$ is a variant of the differentiation instruments proposed in Gandhi and Houde (2019). This instrument measures the magnitude of cost shocks facing a firearm j, relative to cost shocks of rival products j' in the same class. Such variation shifts the pricing behavior and conditional-on-class market shares across firearms within the same class, which serves to identify the lower-level nesting parameter ρ_1 (Verboven 1996).

Our final pair of constraints in Equation (11) uses variation in the size of the choice set, which shifts the class-level market shares, to identify the outside option nesting parameter ρ_0 (Miller and Weinberg 2017, Gandhi and Nevo 2021). We enforce that the class-year preference shocks $\eta_{c,t}$ are uncorrelated with variation in the (de-meaned) count of firearm models $\tilde{J}_{c,t}$ in each class-year, which serves as an additional instrument. Figure D.1 visualizes the year-to-year variation in the size of consumer choice sets, which is driven in large part by firm entry and exit in the firearm industry, such as Remington's bankruptcy in 2020.

Our only remaining demand parameters are the state-year vertical taste shocks $\tau_{s,t}$, which control the generic taste for firearms relative to Massachusetts. We estimate these using measures of firearm purchasing across other state-years. Specifically, and analogous to Equation (9), we estimate $\tau_{s,t}$ by satisfying a set of constraints

$$0 = Pr_{s,t}(j \neq 0|\tau_{s,t}) - \hat{s}_{1,s,t} \quad \forall t, s \neq \text{MA} \tag{12}$$

requiring that the inside share predicted by our model aligns with the inside share observed in NICS data $\hat{s}_{1,s,t}$ each state-year. Thus, our estimate of $\tau_{s,t}$ captures the residual differences in firearm purchasing not explainable by differences in demographics across states. Our approach, identifying preference heterogeneity within a single market with rich microdata, then extrapolating to other markets via aggregate moments, is similar in spirit to the approach taken by Elliott et al. (2023) to infer demand for mobile phones in France.

With a fully estimated demand system, we then recover the unobserved marginal costs $c_{j,t}$ of newly produced firearms using the first order conditions implied by profit maximization in Equation (8).

5.4 Results

Table D.1 presents the full estimates of model parameters governing consumer-specific preference heterogeneity. These estimates imply meaningful heterogeneity in a consumer's willingness-to-pay for different firearms, following intuitive patterns across their observable characteristics. For example, those in more conservative neighborhoods (60% Republican vote share) are willing to pay \$150 more to purchase any firearm $j \neq 0$, and a further \$90 more for a long gun $c(j) = l$, relative to those in less conservative neighborhoods (40% Republican vote share).[35] Figure D.3 shows that the aggregate state-year taste shocks $\tau_{s,t}$ also generate considerable variation in firearm demand across consumers in different states.

Panel (a) of Figure 1 combines both geographic and demographic variation in preferences to show the consumer surplus created by the firearm industry. At its status quo equilibrium each year, the firearms industry creates \$126 of surplus for the average consumer. Across the U.S., this produces \$24.15 billion of consumer surplus each year.[36] Consumer surplus varies considerably across states, with those in the South and Mountain-West tending to

[35] A 20 percentage point change in Republican vote share generates a $2 \times \$76.17 \approx \$150$ increase in marginal willingness-to-pay for handguns relative to the outside option (Table D.2, row 1, column 5). Comparing across consumers in neighborhoods with 60% Republican vote share to those in neighborhoods with 40% Republican vote share implies roughly a $\$70 - -\$20 = \$90$ increase in overall willingness-to-pay for long guns relative to handguns (Figure D.2).

[36] Consumer surplus per U.S. adult is approximately $\$24.15B/255M \approx \$94$, as some adults are not in the market for a firearm.

derive more surplus than along the coasts.

Panel (b) shows that consumers who derive more surplus from the firearms industry are also disproportionately likely to live in congressional districts with many large donations to the NRA. Although excluded from our demand model, NRA donations alone can explain 18% of variation in average consumer surplus across congressional districts (i.e., the R^2 from this univariate linear regression is 0.18). Appendix Figure D.4 shows that these consumers are also more likely to live in areas where their U.S. House Representative is highly rated by the NRA, even among Republican House Representatives. Given the link between the share of pro-gun constituents and the voting records of their representatives (Bouton et al. 2021), these patterns suggest that consumer surplus and the political feasibility of gun regulation are closely related, motivating our consumer surplus constraint on the optimal taxes in Section 2.

Appendix Figure D.5 unpacks substitution patterns across firearms. Own price elasticities range from -1 to -4, with an average own-price elasticity of firearm demand of around -2.6. This is slightly more elastic than the aggregate two-stage least-squares estimate from Section 4, consistent with exogenous cost shocks that affect multiple products within a manufacturer. Our estimated diversion ratios imply that handguns and long guns are poor substitutes for one another, that consumers somewhat substitute between new and used firearms, and that a meaningful share of consumers are willing to substitute to the no-purchase outside option if their preferred firearm raises its price.

Turning to supply, Appendix Figure D.6 presents estimates of margins and marginal cost determinants of firearm manufacturers. We find that the typical firearm has a post-tax Lerner index $(p_{j,t}(1-v_j)-c_{j,t})/p_{j,t}$ between 0.3–0.4. Our model also implies that about $3.5 billion is generated each year in economic profits from firearm sales, similar in magnitude to the $2 billion estimate of accounting profits during our sample period (Khaustovich 2025). Our estimates also imply that firearms with higher capacity, higher caliber, and longer barrel length are more expensive to produce.

In Appendix D.2, we show a close alignment between the predictions of our model and several out-of-sample moments of the national firearms industry, summarized in Figure D.7. Our model captures the time-series of tax revenue generated by the firearms industry, the distribution of production across manufacturers, and cross-state variation in class-specific purchase rates. It also produces estimates of state-year shocks that are orthogonal to conservative vote shares, suggesting that the demographic preferences we estimate in Massachusetts extrapolate to other states.

Summarizing the welfare-relevant lessons of our model, Figure 2 decomposes the surplus from the purchase of an average handgun and long gun model across industry participants.

Firms price the typical handgun at $831. Consumers of the typical handgun are willing to pay $1,063, such that consumers receive $CS_j/(WTP_j - c_j) = 232/(831 + 232 - 436) \approx 37\%$ of the potential surplus from the purchase of the typical handgun. Firms pay $436 of cost to sell the typical handgun, implying that such a sale generates $312 of profit and $83 of tax revenue, or 50% and 13% of potential surplus, respectively. Prices and costs are somewhat higher for long guns, though they generate similar divisions of potential surplus among industry participants.[37]

These divisions of surplus consider only the short-term implications of firearm purchase among industry participants. The following section extends our model to account for the implications of firearm purchase on aggregate homicides over the short and long run.

6 A Model of the Firearm Stock and Homicides

This section specifies a stylized model connecting the flow of firearms in a market to their public health consequences, via downstream firearm homicides. Although stylized, our model accounts for three key forces connecting firearms to homicides: (i) firearm homicides are a function of firearm prevalence (Cook 2018), measured as the stock of firearms in a market; (ii) firearms are durable goods that persist over time; and (iii) firearms affect downstream homicide outcomes differentially, depending on their characteristics. Rather than attempting to estimate the causal effects of firearm purchases on homicides directly, we calibrate our model using existing estimates from the literature.

6.1 Model structure

We specify a two-equation model linking the quantity demanded of new firearms $q^N_{c,s,t}$ by class c in state s during year t, to the contemporaneous firearm stock $Q_{c,s,t}$ and corresponding firearm homicide flow $d_{c,s,t}$:

$$\begin{aligned} \text{(Gun stock law of motion)} \qquad & Q_{c,s,t} = (1-\varphi)Q_{c,s,t-1} + q^N_{c,s,t} \\ \text{(Distribution of gun homicides)} \qquad & d_{c,s,t} \sim \text{Poisson}(Q^{\kappa}_{c,s,t} \cdot \chi_c \cdot \zeta_{s,t}), \end{aligned}$$

independently across class, state, and year. Appendix E provides a microfoundation for this relationship, based on the aggregation of individual decisions to perpetrate interpersonal firearm violence and the risk with which shootings turn fatal.

[37] Figure D.8 shows a similar distribution of surplus from the sales-weighted average handgun and long gun purchase, though these sales-weighted averages transact at lower prices.

The first equation of our model specifies the law of motion for the firearm stock by class, $Q_{c,s,t}$, accounting for the durable nature of firearms. Each year t, a uniformly random fraction $\varphi \in (0,1)$ of firearms within a state-class are scrapped and exit circulation, while a fraction $1-\varphi$ persist in the stock. Then, an additional flow of new firearms $q^N_{c,s,t} = \sum_{j:c(j)=c, j\in\mathcal{J}^N_t} q_{j,s,t}$ from the contemporaneous legal firearms market is added to the stock. This law of motion captures that, even if the U.S. were to ban the sales of new firearms altogether, firearm ownership would persist, due to the large stock of existing firearms. We treat the existing firearm stock in 2015 $Q_{c,s,2015}$ as an initial condition to be calibrated from data, and use predictions from our model of the firearm market to determine the stock and flow of new firearms $q^N_{c,s,t}$ each year from 2016-2022.[38]

The second equation models firearm homicides $d_{c,s,t}$ as a Poisson random variable with expectation $Q^{\kappa}_{c,s,t} \cdot \chi_c \cdot \zeta_{s,t}$. This functional form conforms to the literature, in specifying a constant-elasticity relationship κ between firearm homicides $d_{c,s,t}$ and the contemporaneous firearm stock $Q_{c,s,t}$, undifferentiated across class c (Duggan 2001, Azrael et al. 2004, Cook and Ludwig 2006, Kim and Wilbur 2022). The class-specific scalars χ_h, χ_l control the scale of the relationship between the firearm stock and homicides within each class, with χ_h/χ_l measuring the relative "efficiency" of handguns against long guns in the production of total firearm homicides $d_{s,t} = d_{h,s,t} + d_{l,s,t}$. The remaining term $\zeta_{s,t}$ is an unobserved state-year firearm homicide shock, capturing other determinants of expected firearm homicides across markets (e.g., differences in criminal justice policy) that are constant across class c, but may be arbitrarily correlated with the stock $Q_{c,s,t-1}$ and flow $q^N_{c,s,t}$ of firearms.

The above model provides a useful mapping between the legal firearm market and the occurrence of firearm homicides. Specifically, the change in homicides with respect to a new firearm purchase is[39]

$$\frac{\partial E[d_{c,s,t}]}{\partial q^N_{c,s,t}} = \kappa\Big((1-\varphi)Q_{c,s,t-1} + q^N_{c,s,t}\Big)^{\kappa-1} \cdot \chi_c \cdot \zeta_{s,t} \tag{13}$$

where the terms on the right-hand side depend only on the fixed parameters and equilibrium outcomes from our model.

Before turning to estimation, two assumptions of this model are important to highlight. We focus on the transmission of demand in the legal market to firearm homicides, abstracting

[38] We implicitly assume that used guns captured by the composite good ω_c are already counted in the firearms stock $Q_{c,s,t}$. Appendix E provides a robustness check where we allow used gun purchases to affect the stock, and results are qualitatively similar.

[39] To drop the expectation on the left-hand side, scale the right-hand side by the random variable $\hat{e}_{c,s,t}$ to represent the purely random Poisson variability: $d_{c,s,t} = Q^{\kappa}_{c,s,t} \cdot \chi_c \cdot \zeta_{s,t} \cdot \hat{e}_{c,s,t}$, with $E[\hat{e}_{c,s,t}] = 1$ and $Pr(\hat{e}_{c,s,t} \geq 0) = 1$.

away from behavior on the illicit market (Cook et al. 2007). While the perpetrators of gun crime often illegally acquire their firearms, the vast majority of guns used in crime originate in the legal market.[40] Moreover, the homicide elasticities we use for calibration are estimated from proxies for overall gun prevalence, so they capture both the direct effect of legally owned guns on homicides and the indirect effects from diversion to the illicit market.

An additional assumption, we model homicides as a function of the state-specific firearm stock. While some guns used in crime cross state lines, most are locally sourced. During our sample period, 72% of guns recovered from crimes and traced were purchased from a dealer in the state where the crime occurred (ATF 2023).

6.2 Calibration

We calibrate the parameters in Equation (13) using measurements from the existing literature, with calibration targets and parameter estimates summarized in Table 4. Appendix E provides a detailed explanation of the data and moments we use for calibration, as well as our procedures for inference.

We calibrate the initial firearm stock $Q_{c,s,2015}$ by decomposing it into separate quantities, which can be individually measured from information in the 2015 ACS, proxies for firearm ownership per capita from Schell et al. (2020), and microdata from the 2015 National Firearm Survey (Azrael et al. 2017).

To calibrate the scrap rate φ of firearms from the stock, we align existing survey estimates of the national firearm stock at two distinct points in time Q_{1994}, Q_{2015} with the ATF's measure of the inflow of new firearms over the intervening periods q_t^N and our assumed law of motion:

$$Q_{2015} = (1-\varphi)^{2015-1994} Q_{1994} + \sum_{t=1995}^{2015} (1-\varphi)^{2015-t} q_t^N.$$

Solving this equation leads to a calibrated scrap rate of $\varphi = 0.015$. Although less-likely to be scrapped than other consumer durables (Jacobsen and Van Benthem 2015), our calibrated scrap rate for firearms φ is higher than measured or assumed in prior studies of the firearm market (Moody 2010, Azrael et al. 2017, McDougal et al. 2023).

Turning to health, our calibration captures the average relationship between firearms and gun homicides by matching existing estimates of the elasticity of firearm homicides with

[40] About 97% of guns recovered from crime scenes and traced by the ATF were originally purchased in the legal market (ATF 2023). Sources outside the legal market are privately made firearms, which are never registered (2.0%), and the guns reported as thefts or losses by FLLS, which are never sold (0.8%).

respect to the overall firearm stock, approximately 0.294.[41] As these existing estimates do not distinguish between firearms of different classes, we use this estimate to calibrate the aggregate elasticity of our model with respect to a change in the stock of each firearm class:

$$0.294 = \hat{E}\left[\frac{\partial \log d_{s,t}}{\partial \log Q_{s,t}}\right] \overset{\text{calibrate}}{=} E\left[\sum_c \frac{\partial \log d_{s,t}}{\partial \log Q_{c,s,t}}\right] = \kappa.$$

Appendix E gives a sufficient condition for this calibration: that the empirical variation in the firearm stock $Q_{s,t}$ used to estimate the left-hand elasticity does not alter the conditional share of firearms within each class $Q_{c,s,t}/Q_{s,t}$.

To calibrate the relative homicide "efficiency" χ_h/χ_l of handguns to long guns, we match our model's predicted share of handguns among firearm homicides in the average state-year to the value observed in CDC vital statistics:[42]

$$0.899 = \sum_{s,t} \frac{d_{s,t}}{\sum_{s',t'} d_{s',t'}} \times \frac{d_{h,s,t}}{d_{s,t}} \overset{\text{calibrate}}{=} \sum_{s,t} \frac{d_{s,t}}{\sum_{s',t'} d_{s',t'}} \times \frac{\frac{\chi_h}{\chi_l} Q_{h,s,t}^{\kappa}}{\frac{\chi_h}{\chi_l} Q_{h,s,t}^{\kappa} + Q_{l,s,t}^{\kappa}}.$$

This procedure leads to a value of $\chi_h/\chi_l = 8.9$, which implies that a handgun purchase produces roughly nine times more homicides than a long gun purchase, holding equal the stock of firearms in each class.

6.3 Results

Panel D of Table 4 shows the expected change in contemporaneous homicides from the sale of each new firearm in our sample. We compute these quantities by evaluating Equation (13) purchase-by-purchase, averaging over the distribution of new firearm purchases $q_{j,s,t}^N$ predicted by the model. The average handgun sale in the U.S. from 2016–2022 caused 3.35×10^{-5} firearm homicides during the year in which it was purchased. The average long gun sale over the same period caused 3.61×10^{-6} firearm homicides. Since handguns accounted for around 75% of new firearm sales over this period, the average firearm purchase in the U.S. caused 2.63×10^{-5} homicides during its first year in the firearm stock.

We find that firearms are scrapped from the stock at rate $\varphi = 0.015$ per year, so that

[41] Cook and Ludwig (2006) estimates an elasticity of 0.272, using the fraction of suicides committed by guns as a proxy measure. Duggan (2001) estimates a very similar elasticity of 0.316 using subscriptions to *Guns and Ammo* as a proxy measure for gun prevalence. We take the unweighted average of these two estimates.

[42] The left-hand side of this equation is a convex combination of the handgun share of firearm homicides each state-year $d_{h,s,t}/d_{s,t}$, observed in the National Violent Death Reporting System. Where the weights are equal to each state-year's share of all firearm homicides in the U.S. during our sample, observed in the Multiple Cause of Death data from the National Center for Health Statistics.

the typical new firearm remains in circulation for $1/\varphi \approx 67$ years. Thus, our model implies that the average new firearm purchase in the U.S. from 2016–2022 would generate ≈ 0.00175 homicides over its lifetime, or around 1 homicide for every 570 new firearm purchases.

We convert the lifetime homicide flows from a new firearm purchase into dollars of net present expected value under two further calibrations. Specifically, we impose a yearly discount factor of 0.95 and a statistical value of a life equal to \$9.5M (Peterson et al. 2021). Under these assumptions, the average new firearm purchased between 2016–2022 generates \$3,868 of present-day harm from its lifetime firearm homicide flows.

Figure 2 compares the homicide costs of firearm purchase between handguns and long guns, and to the surplus generated for industry participants. The typical sale of one more handgun from 2016–2022 would generate around \$315 of damage from firearm homicides in its first year of circulation. Although smaller in magnitude than market surplus, this one-year homicide cost is 36% higher than the surplus to the consumer from purchasing a handgun. However, the market surplus (taxes+profits+CS) of \$627 from the typical new handgun purchase is dwarfed by its net present value of expected homicide fatalities, equal to \$4,881, accounting for the durability of firearms. Notably, the homicide costs of a new long gun purchase are about 10% of the homicide costs attributable to handguns, which is sufficiently low that each long gun sale produces positive total surplus. These patterns speak to the ability of regulators to improve welfare with the design of targeted firearm regulation, especially through shifting tax rates towards handguns, which we explore in the following section.

We consider the robustness of our model of firearm homicides in Appendix E, as summarized by Table E.1. These robustness checks are important, given the stylized nature of the model in this Section. We evaluate departures from our baseline modeling assumptions: allowing separate homicide elasticities for handguns κ_h and long guns κ_l (Cook and Ludwig 2006, Braga and Cook 2018); assuming that the convenience of carrying handguns as an input to crime has no role in their homicide "efficiency" χ_h/χ_l relative to long guns (Libby and Corzine 2007); assigning zero social cost to the subset of justifiable homicides committed by private citizens; adding the flow of used firearms ω_c as an input to the homicide process; and apportioning harms from a consumer's purchase of a new firearm, regardless of their choice of firearm class c, using only their immutable preference characteristics $(\alpha_i, \beta_i, \vec{\epsilon}, \delta_{s,j,t})$. Each of these alternatives preserves our baseline model's key qualitative implications and lessons for policy design, as shown in Figure E.1.

7 Counterfactuals

This section combines our fitted models of the firearm industry and firearm homicides from Sections 5 and 6—embedding them into the optimal tax model of Section 2—to study the design of alternate firearm regulation in the U.S.

7.1 Policies, Implementation, and Interpretation

We consider counterfactual changes in the federal excise tax on the sales of new firearms from manufacturers v_j. Each counterfactual can be represented by the vector of tax changes $\Delta\vec{v}$ of length $\sum_{t=2016}^{2022} |\mathcal{J}_t^N|$. The status quo of a 10% tax on handguns and an 11% tax on long guns is implemented by $\Delta\vec{v} = \vec{0}$. We benchmark the political feasibility constraint at consumer surplus under the status quo $CS_0 = CS(\Delta\vec{v} = \vec{0})$.

Our analysis implements three progressively complex counterfactual tax changes $\Delta\vec{v}$. The first is a homogeneous tax increase of 11 percentage points across all firearms $\Delta\vec{v} = 0.11$, such that the tax rate on handguns is 21% and on long guns is 22%. This counterfactual scales up state-level firearm tax policies to the whole of the U.S. (Brownlee 2024), particularly California AB 28, which implemented an identical tax increase in 2024 in California.[43] Due to its real-world policy relevance, we evaluate both the welfare effects of this tax change and its political feasibility $CS(\Delta\vec{v} = 0.11) - CS_0$, as formalized in Section 2.

We also consider a simplified implementation of the politically constrained optimal taxes of Section 2, allowing only a "tilt" of the status-quo tax schedule for handguns and long guns. Specifically, we constrain the tax scheme to vary only by weapon class, so that all firearms in the same class c face the same tax change. Since we do not develop theoretical results when tax instruments are constrained, we implement this counterfactual by numerically solving the constrained optimization in Equation (2) over the two-dimensional space of class-specific tax changes $(\Delta v_h, \Delta v_l)$. In solving this problem, we also impose that neither tax change can lead to a net subsidy, i.e. $v_c + \Delta v_c \geq 0$. Empirically, our solution removes the tax on the low-externality class (long guns, $v_l + \Delta v_l = 0$), while setting the uniformly highest rate on the high-externality class (handguns, $v_h + \Delta v_h = 15.5\%$) that holds consumer surplus at the status quo CS_0. Since this tax system sets no tax on long guns and a uniform tax on handguns, we refer to it as the uniform handgun tax.

Our most sophisticated counterfactual approximates the full solution of the planner's problem in Equation (2), potentially requiring a distinct tax change for every firearm $\Delta\vec{v}$. We compute these counterfactual taxes by plugging in values for the objects in Equation (3) using estimates from our counterfactual equilibria, following a sufficient statistics approach

[43] See this link for the legislation: https://legiscan.com/CA/text/AB28/id/2842856.

for estimating this tax (Allcott et al. 2019). The taxes we compute only approximate the full solution, because Equation (3) may imply different taxes at different equilibria (i.e. the inputs to the tax are themselves equilibrium objects). Appendix F provides further details on the implementation.

In practice, this tax system imposes no tax on long guns, but heterogeneous rates on each handgun each year, so we refer to it as the targeted handgun tax.

For each counterfactual, we alter the status quo tax rates on new firearm sales from 2016–2022 by $\Delta\vec{v}$, and trace its effects on firearm pricing, purchase, ownership, and homicides during our sample. Prices adjust, as firearm producers re-optimize to maximize profits under the tax change $\Delta\vec{v}$, facing the same demand system, ownership matrix, marginal costs, used-firearm values $\omega_{c,s,t}$, and Nash-Bertrand pricing conditions we recover in Section 5. Facing different prices, consumers adjust their firearm purchases, which alters the profits, federal tax revenues, and consumer surplus from the industry. The change in firearm purchases also affects firearm homicides—as in Equation (13)—because the tax change $\Delta\vec{v}$ adjusts the endogenous evolution of the firearm stock.

We measure the change in welfare from a counterfactual tax change $\Delta\mathcal{W}(\Delta\vec{v})$ by its component effects on consumer surplus, profits, federal government revenue, and homicide externalities, as in Equation (1). For each counterfactual, we report the change in welfare relative to the status quo from 2016–2022:

$$\Delta\mathcal{W}(\Delta\vec{v}) \equiv \sum_{t=2016}^{2022} \mathcal{W}_t(\Delta\vec{v}) - \mathcal{W}_t(\vec{0}),$$

where the change in each welfare component is defined analogously. Appendix F provides explicit expressions.

Unlike our analysis of the status quo equilibrium in Sections 3–6, our analysis of welfare under counterfactual taxes $\Delta\vec{v}$ involves both extrapolating beyond the observed data and assigning normative interpretations to our model, both of which limit the reliability of our results.

In extrapolating beyond the data, our counterfactuals allow only adjustments to the prices and quantities of legal firearms already available to consumers, holding fixed large components of the firearms industry. Specifically, we do not allow for the possibility that a tax change may affect the innovation of new firearm models, the equilibrium of the illicit

firearm market,[44] the behavior of firearm retailers, or the valuations of used firearms.[45] Our assumption that the value of used firearms $u_{i,\omega_c,s,t}$ is invariant to the tax change $\Delta\vec{v}$ is a substantive restriction on both supply and demand.[46] This assumption is satisfied if used firearms are supplied by a perfectly elastic competitive fringe, so that their price is constant across variation in the underlying firearm stock and changes in consumer demand.[47]

Among the normative interpretations we impose, several key ones are that only federal tax revenues enter welfare,[48] that dollars of tax revenue are exchangeable with dollars of market surplus,[49] that marginal costs of production are a pure welfare loss,[50] and that the only externality from firearm purchase is the statistical value of firearm homicides.[51] We view these caveats as necessary to appropriately interpret our predictions of the effects of counterfactual tax changes $\Delta\vec{v}$. Though, our choice of counterfactuals that hold consumer surplus constant limits the degree of extrapolation from the current equilibrium.

7.2 Counterfactual Results

Panel (a) of Figure 3 plots the average annual change in welfare from 2016–2022 for each of the counterfactual taxes $\Delta\vec{v}$ we consider. Our simplest counterfactual—raising the existing tax rate by 11 percentage points—improves overall welfare by 280M$ per year during our sample. These gains are achieved in roughly equal proportion through an $800M increase

[44] Existing work suggests minimal consumer substitution between the licit and illicit firearm markets (Cook et al. 2007), with comparable firearms transacting at 300% higher prices through illicit channels. The price changes we study in our counterfactuals are smaller by an order of magnitude, giving us confidence in our extrapolation.

[45] We match the structure of current federal firearm tax policies by assessing counterfactual taxes only on the sale of new (not used) firearms.

[46] It restricts supply, for example, because manufacturers do not internalize that their sale of a new firearm in period t may increase the supply of, and competition from, used firearms in a future period t'. It restricts demand, for example, because consumers do not internalize that a newly purchased firearm can be resold tax-free.

[47] We view this as a reasonable approximation of the U.S. during our sample period. Used firearms could be supplied from more than 300 million firearms in circulation, and they could be bought and sold at tens of thousands of licensed, brick-and-mortar suppliers (Azrael et al. 2017, Johnson and Robinson 2021). Our counterfactual taxes change the total flow of guns by less than 0.5 million units per year, so even an upward sloping supply curve would have a negligible impact on prices.

[48] Firearm purchases also generate tax revenue for state and local governments through sales taxes and licensing fees, absent from our model.

[49] This abstracts from both the marginal value and cost of public funds (O'Connell and Smith 2024). Federal firearm taxes are currently earmarked for wildlife conservation (Congressional Research Service 2023), which may generate economies of scale and/or intrinsic, un-capitalized value.

[50] Marginal costs may include payments to workers (e.g., wages), to upstream suppliers (e.g., materials), or to downstream distributors (e.g., marketing), all of which may also generate social welfare.

[51] Firearms are also used in non-lethal crime. Though, past studies have found smaller effects of the firearm stock on non-fatal crimes (Duggan 2001, Rosenberg 2025). We also do not account for any effects of the firearm stock on incidents of self-harm, nor on broader externalities like social well-being, mental health, or schooling (Pienkny et al. 2024, Cabral et al. 2021).

in federal tax revenue and the prevention of 80 firearm homicides each year. However, this homogeneous tax increase on all new firearms considerably harms direct market participants, destroying $659M of yearly consumer surplus and $618M of yearly manufacturer profits.

In contrast, our uniform handgun tax achieves a 41% higher overall welfare gain. This system is more efficient because it has an independent policy instrument for each firearm class, allowing it to better align each firearm's tax rate with its social cost, especially as Figure 2 documents that handguns are an order of magnitude more harmful than long guns. To achieve this efficiency gain, the uniform handgun tax prevents 36 firearm homicides each year, while essentially holding market participants' surplus and tax revenues constant. By maintaining surplus among many stakeholders, while targeting only consumer surplus, these results suggest that our narrow, stylized constraint on political feasibility may also support the design of policies that generate broader political support (e.g. among firms).

The targeted handgun tax generates a further 26% increase in welfare relative to the uniform handgun tax. Panel (a) of Figure 4 shows that this welfare gain comes from further differentiating tax rates across handguns, with meaningful shares of handguns facing rates from 0%–26%. Panels (b)-(d) provide information on the sort of handgun models that are taxed more: those that have lower production costs ("Saturday night specials"), provide less value to consumers, and/or are more substitutable with less-harmful alternatives (e.g. the outside option). Heuristically, the constrained-optimal firearm taxes subsidize the inframarginal consumers of new handguns, in order to target higher taxes at consumers who may be diverted away from new handguns on the margin. Through this mechanism, the targeted handgun tax reduces the flow of new firearms into circulation, preventing 47 firearm homicides each year, while also increasing tax revenue at the cost of some manufacturer profits.

The welfare effects of these counterfactual firearm taxes arise from changes in the equilibrium behavior of supply and demand, summarized in Appendix Figure F.1. Under our maintained assumption of Nash-Bertrand competition, market power among manufacturers leads them to incompletely pass-through firearm taxes, with only about 60% of the tax change $\Delta\vec{v}$ converting to consumer prices. At equilibrium prices, raising the firearm tax by 11 percentage points would reduce the total quantity of firearm purchases by 360,000 units per year. In contrast, our politically constrained counterfactual taxes would reduce quantity by less than 65,000 units per year (80% less), and instead shift the composition of firearm purchases away from new handguns and towards new long guns and used handguns. The targeted handgun tax, inducing greater equilibrium price dispersion, further shifts the composition away from new handguns. Figure F.2 shows that erroneously assuming perfect competition when evaluating these counterfactual tax systems would overstate their welfare

effects by an order of magnitude and incorrectly identify the most efficient policy.[52]

We also consider regulation through a change in the structure of firearm supply, with Figure F.3 evaluating counterfactual outcomes under the static, multi-product monopoly markups that could obtain under full price coordination. This counterfactual is motivated both by theory—Buchanan (1969) argues that monopoly power may correct externalities without regulatory intervention—and by the legal structure of the U.S. Master Settlement Agreement in the tobacco industry (Cutler et al. 2002). Relative to the counterfactual taxes we consider, full monopolization of firearms would generate larger changes on every component of social welfare, and larger overall welfare gains. Full monopolization of firearms also meets a key feasibility condition of the tobacco settlement, generating enough revenue that its division via lump-sum transfer could increase manufacturer profits and federal tax revenues above the status quo (Bulow and Klemperer 1998). However, this condition may be insufficient to support regulation through monopolization in the firearms industry, where consumers and their advocacy groups wield considerable political power, and full monopolization would destroy 10% of the status quo consumer surplus.

To benchmark the effects of these counterfactual regulations during our sample period, we compare their public health consequences to existing firearm policies that have been the subject of prior research, as compiled by Smart et al. (2024). Table F.1 shows that the politically feasible counterfactual taxes implement 3%–7% of the public health gains of many existing policies, like implementing universal background checks, while limiting economic harm to market participants. Appendix F.3 provides details.

Since the typical firearm remains in circulation for 67 years, considering only the period 2016–2022 may mask longer term benefits of a firearm tax change. Without extrapolating beyond our sample, Panel (b) of Figure 3 plots the annual accumulated change in social welfare under each counterfactual tax. Each tax improves over time, as it continually diminishes the stock of socially harmful firearms. After 7 years, simply raising the tax rate on all firearms by 11 percentage points is the best-performing policy, highlighting that higher firearm taxes may produce substantial welfare gains through homicide reduction over the long-run. These gains accrue slowly, while losses in market surplus are immediate, and this homogeneous tax increase lowers welfare below the status quo during its first two years of implementation. Conversely, there are immediate welfare gains from the firearm taxes that meet our stylized constraint on political feasibility, suggesting this design may also generate political support within the short-term electoral cycles that often motivate policy (Bonfiglioli

[52]Explicitly, we assume firms set $p_{j,t}(1-v_{j,t}) = c_{j,t}$ and make zero economic profits. Under this assumption, we re-solve for the uniform and targeted handgun taxes that set consumer surplus equal to zero, using the optimal targeted tax from Corollary 1 in Appendix A.

and Gancia 2013).

7.3 The Feasibility of Counterfactual Firearm Taxes

We consider notions of feasibility for firearm taxes that extend beyond the stylized political constraint in Section 2.

Although we motivated political feasibility as a constraint on consumer surplus—by discussing the influence of firearm owners and their advocacy organizations—our empirical analysis reveals that certain subgroups of consumers may be especially pivotal for policy design. National surveys report that firearm ownership is two times higher among Republicans than Democrats (Parker et al. 2017), we observe more firearm purchasing among consumers in conservative parts of Massachusetts, and our fitted model reveals that consumers enjoy more surplus from the firearm market in parts of the U.S. with higher NRA donations and Republican vote shares. Insofar as changing firearm taxes at the federal level would require support from Republican legislators, their decisions may reflect the effects of the tax on their constituents and political base.

Figure 5 examines political feasibility by directly plotting the welfare effects of each counterfactual tax against the Republican vote share in the 2016 Presidential Election, state-by-state. Our measure of state-level welfare ("population surplus") includes only consumer surplus and prevented firearm homicides, to represent that these benefits accrue directly to the state population.

Panel (a) shows how a homogeneous increase in the federal firearm tax would affect each state. This tax destroys considerable consumer surplus, disproportionately in states with higher Republican vote share, while gains in public health are diffuse across the political spectrum. As such, the change in population surplus is inversely correlated with Republican vote share. California, where the state government implemented such a tax in 2024, would enjoy a net benefit from this policy change. However, 29 states would experience a net loss in population surplus, underscoring the political challenges of passing a broad scale increase in firearm taxes at the national level.

Panels (b) and (c) plot the effects of our two politically-constrained taxes, which induce similar profiles. National consumer surplus remains at the status quo, by design, but the taxes satisfying this constraint also slightly shift consumer surplus towards states with a higher Republican vote share. Consumers in conservative areas have a disproportionately higher taste for long guns (with counterfactually zero tax) than for handguns (with counterfactually higher tax), as demonstrated in Appendix Figure D.2. This preference heterogeneity also produces larger public health gains in more Republican states, as purchases by

consumers with a disproportionately low taste for handguns can be more efficiently diverted away from new handguns on the margin. Together, these forces reduce firearm homicides with no systematic loss in consumer surplus, generating a net gain in population surplus for every state, no matter its position on the political spectrum.

Beyond examining the feasibility of counterfactual taxes under our stylized political constraint, and across the political spectrum of U.S states, we consider feasibility among narrower subsets of consumers. Appendix Figure F.4 shows that the taxes designed under our political constraint also shift consumer surplus into congressional districts with more NRA activity.

At a fundamental level, our counterfactuals subsidize long guns to tax handguns, which we believe does not preclude their political feasibility (i.e., there is not a "blocking coalition" of consumers who prefer handguns). Consumers with a higher taste for handguns tend to live in areas that are more urban and have more liberal political views, both characteristics of gun owners that support tighter firearm regulation (Parker et al. 2017). Moreover, 82% of handgun owners in the NRA also own a long gun, as do 60% of handgun owners nationwide (Berrigan et al. 2022),[53] such that the relative subsidies and taxes we analyze can average out within an individual across their separate firearm purchase occasions.[54]

Across our analysis of counterfactuals, the uniform handgun tax appears as a potentially strong policy option. It maintains consumer surplus, manufacturer profits, and federal revenues at the status quo, while saving dozens of lives each year. These welfare gains arrive quickly and grow in magnitude over time. By only adjusting two tax rates that already exist in federal law, this system is simple to explain and to codify. Yet, it achieves around 80% of the welfare gain of our targeted tax system, which uses theory to calibrate a different tax rate for each firearm, each year. Perhaps most importantly, the uniform handgun tax does no systematic harm to consumers across the political spectrum, while providing disproportionate benefits to public health in Republican- and NRA-supporting areas. As such, we view the uniform handgun tax as a feasible, politically palatable improvement over the existing federal taxes on the consumer firearms industry.

[53] Source: authors' own tabulations from the survey microdata in Berrigan et al. (2022).

[54] Our demand model of Section 5 is compatible with these facts. The nesting structure on the idiosyncratic utility component $\epsilon_{i,j,s,t}$ leads a consumer to have minimal substitution between firearm classes each year, while i.i.d. draws of $\epsilon_{i,j,s,t}$ across years ("occasions") lead that same consumer to purchase both handguns and long guns over time.

8 Conclusion

In this paper, we provide estimates of supply and demand for the U.S. consumer firearm industry that are based on observed firearm transactions and the prices facing consumers. Through descriptive analysis and a structural model, we show that firearm consumers are moderately price elastic, that demographics are a substantial determinant of consumer preferences for firearms with different characteristics, and that substitution across firearm classes is relatively low. We estimate that market power plays a significant role in the supply side of the firearms market, leading to large markups and incomplete pass-through of taxes to consumers, which meaningfully impacts the welfare implications of alternative policies. In order to benchmark the effect of firearm market outcomes on downstream public health externalities, we calibrate a stylized model of the firearm stock and its effect on firearm homicides. In our counterfactual policy exercises, we show that targeted taxes on handguns are more welfare enhancing than uniform tax increases, particularly in the short-to-medium run, primarily because purchases of handguns generate more homicides than long guns. Moreover, these targeted, handgun-only policies have properties that may make them politically feasible: gun consumers lose no surplus on net, and the largest beneficiaries are those living in conservative areas, whose constituents have historically been a barrier to new firearm legislation.

While our paper is a step towards understanding this market, it leaves many questions unanswered. This paper does not speak to behavior on the illicit market for firearms, and how this may respond to changes in government policy. Nor does this paper explicitly model dynamics in the gun market, though we show that dynamics are important for understanding the public health benefits of regulation, because guns are highly durable goods. Finally, this paper cannot speak to larger-scale reforms of firearms regulation, such as universal buyback programs or bans on certain types of firearms Addressing these questions is challenging, especially due to the limitations facing researchers on assembling data about the firearms industry. Our hope is that further work on this topic can aid both researchers' understanding of this market and policymakers ability to design effective, implementable regulations for the consumer firearm industry.

References

Hunt Allcott, Benjamin B Lockwood, and Dmitry Taubinsky. Regressive sin taxes, with an application to the optimal soda tax. *The Quarterly Journal of Economics*, 134(3): 1557–1626, 2019.

Dmitry Arkhangelsky, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager. Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118, 2021.

Luis Armona and Adam M Rosenberg. Measuring the market for legal firearms. In *AEA Papers and Proceedings*, volume 114, pages 52–57. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2024.

John Asker, Allan Collard-Wexler, Charlotte De Canniere, Jan De Loecker, and Christopher R Knittel. Two wrongs can sometimes make a right: The environmental benefits of market power in oil. Technical report, National Bureau of Economic Research, 2024.

ATF. National firearms commerce and trafficking assessment (NFCTA): Crime guns – volume two. Technical report, U.S. Department of Justice, Washington, DC, February 2023. URL https://www.atf.gov/firearms/national-firearms-commerce-and-trafficking-assessment-nfcta-crime-guns-volume-two. Accessed August 19, 2026.

Deborah Azrael, Philip J Cook, and Matthew Miller. State and local prevalence of firearms ownership measurement, structure, and trends. *Journal of Quantitative Criminology*, 20: 43–62, 2004.

Deborah Azrael, Lisa Hepburn, David Hemenway, and Matthew Miller. The stock and flow of us firearms: results from the 2015 national firearms survey. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 3(5):38–57, 2017.

Meenakshi Balakrishna and Kenneth C Wilbur. Do firearm markets comply with firearm restrictions? how the massachusetts assault weapons ban enforcement notice changed registered firearm sales. *Journal of empirical legal studies*, 19(1):60–89, 2022.

Nano Barahona, Cristóbal Otero, and Sebastián Otero. Equilibrium effects of food labeling policies. *Econometrica*, 91(3):839–868, 2023.

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.

John Berrigan, Deborah Azrael, and Matthew Miller. The number and type of private firearms in the united states. *The ANNALS of the American Academy of Political and Social Science*, 704(1):70–90, 2022.

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.

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.

Frederick J Boehmke, Sean Gailmard, and John W Patty. Business as usual: interest group access and representation across policy-making venues. *Journal of Public Policy*, 33(1): 3–33, 2013.

Katie Bollman, Benjamin Hansen, Edward A. Rubin, and Garrett O. Stanford. Gun policy and the steel paradox: Evidence from oregunians. Technical report, National Bureau of Economic Research, 2025.

Alessandra Bonfiglioli and Gino Gancia. Uncertainty, electoral incentives and political myopia. *The Economic Journal*, 123(568):373–400, 2013.

Laurent Bouton, Paola Conconi, Francisco Pino, and Maurizio Zanardi. The tyranny of the single-minded: Guns, environment, and abortion. *Review of economics and statistics*, 103 (1):48–59, 2021.

Anthony A Braga and Philip J Cook. The association of firearm caliber with likelihood of death from gunshot injury in criminal assaults. *JAMA network open*, 1(3):e180833–e180833, 2018.

Anthony A Braga and David M Hureau. Strong gun laws are not enough: the need for improved enforcement of secondhand gun transfer laws in massachusetts. *Preventive medicine*, 79:37–42, 2015.

Anthony A Braga, Elizabeth Griffiths, Keller Sheppard, and Stephen Douglas. Firearm instrumentality: do guns make violent situations more lethal? *Annual Review of Criminology*, 4(1):147–164, 2021.

Jurgen Brauer. The us firearms industry: Production and supply. *Small Arms Survey Working Paper 14*, 2013.

Geoffrey Brennan and James M Buchanan. *The power to tax: Analytic foundations of a fiscal constitution*. Cambridge University Press, 1980.

Chip Brownlee. Seven states move to tax guns and ammo. 2024.

Jan K Brueckner. Airport congestion when carriers have market power. *American Economic Review*, 92(5):1357–1375, 2002.

James M Buchanan. External diseconomies, corrective taxes, and market structure. *The American Economic Review*, 59(1):174–177, 1969.

Jeremy Bulow and Paul Klemperer. The tobacco deal. *Brookings Papers on Economic Activity. Microeconomics*, 1998:323–394, 1998.

Marika Cabral, Bokyung Kim, Maya Rossin-Slater, Molly Schnell, and Hannes Schwandt. Trauma at school: The impacts of shootings on students' human capital and economic outcomes. Technical report, National Bureau of Economic Research, 2021.

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.

William G Cochran. The combination of estimates from different experiments. *Biometrics*, 10(1):101–129, 1954.

Congressional Research Service. Firearms and ammunition excise tax (faet). 2023.

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 Nirupama L Rao. The cost of curbing externalities with market power: Alcohol regulations and tax alternatives. Technical report, National Bureau of Economic Research, 2023.

Philip J Cook. Gun markets. *Annual Review of Criminology*, 1(1):359–377, 2018.

Philip J Cook and Jens Ludwig. *Guns in America: results of a comprehensive national survey on firearms ownership and use*. Police Foundation Washington, DC, 1996.

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

Philip J Cook, Jens Ludwig, Sudhir Venkatesh, and Anthony A Braga. Underground gun markets. *The Economic Journal*, 117(524):F588–F618, 2007.

David M Cutler, Jonathan Gruber, Raymond S Hartman, Mary Beth Landrum, Joseph P Newhouse, and Meredith B Rosenthal. The economic impacts of the tobacco settlement. *Journal of Policy Analysis and Management: The Journal of the Association for Public Policy Analysis and Management*, 21(1):1–19, 2002.

André de Palma and Karim Kilani. Transition choice probabilities and welfare analysis in additive random utility models. *Economic Theory*, 46(3):427–454, 2011.

Sara Drango, Sarah Moshary, and Bradley Shapiro. California's firearm excise tax is almost fully passed on to consumers. Technical report, National Bureau of Economic Research, 2025.

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

Isis Durrmeyer and Mario Samano. To rebate or not to rebate: Fuel economy standards versus feebates. *The Economic Journal*, 128(616):3076–3116, 2018.

Jonathan T Elliott, Georges V Houngbonon, Marc Ivaldi, and Paul T Scott. Market structure, investment, and technical efficiencies in mobile telecommunications. 2023.

Amit Gandhi and Jean-François Houde. Measuring substitution patterns in differentiated-products industries. *NBER Working paper*, (w26375), 2019.

Amit Gandhi and Aviv Nevo. Empirical models of demand and supply in differentiated products industries. In *Handbook of industrial organization*, volume 4, pages 63–139. Elsevier, 2021.

Austan Goolsbee and Amil Petrin. The consumer gains from direct broadcast satellites and the competition with cable tv. *Econometrica*, 72(2):351–381, 2004.

Laura Grigolon and Frank Verboven. Nested logit or random coefficients logit? a comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics*, 96(5):916–935, 2014.

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

Niklas Hüther. More guns lead to more crime: Evidence from private equity deals. *Mimeo*, 2023.

Janice Iwama and Jack McDevitt. Rising gun sales in the wake of mass shootings and gun legislation. *Journal of Primary Prevention*, 42:27–42, 2021.

Mark R Jacobsen and Arthur A Van Benthem. Vehicle scrappage and gasoline policy. *American Economic Review*, 105(3):1312–1338, 2015.

David B. Johnson, Joshua J. Robinson, Daniel Semenaza, and Alexi Thompson. Where are the guns? Technical report, 2023. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4119613.

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

Aaron Karp. Estimating global civilian-held firearms numbers. 2018.

Vlad Khaustovich. Guns and ammunition manufacturing in the us. Technical report, IBISWorld, 2025.

Jessica Jumee Kim and Kenneth C Wilbur. Proxies for legal firearm prevalence. *Quantitative Marketing and Economics*, 20(3):239–273, 2022.

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

Christopher S Koper and Jeffrey A Roth. The impact of the 1994 federal assault weapons ban on gun markets: An assessment of short-term primary and secondary market effects. *Journal of quantitative criminology*, 18(3):239–266, 2002.

Matthew J Lacombe. *Firepower: How the NRA turned gun owners into a political force*. Princeton University Press, 2021.

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

Eulalie Laschever and David S Meyer. Growth and decline of opposing movements: Gun control and gun rights, 1945–2015. *Mobilization*, 26(1):1–20, 2021.

Lois K Lee, Katherine Douglas, and David Hemenway. Crossing lines—a change in the leading cause of death among us children. *New England Journal of Medicine*, 386(16): 1485–1487, 2022.

Nicholas E Libby and Jay Corzine. Lethal weapons: Effects of firearm types on the outcome of violent encounters. *Justice Research and Policy*, 9(2):113–137, 2007.

Richard G Lipsey and Kelvin Lancaster. The general theory of second best. *The review of economic studies*, 24(1):11–32, 1956.

Erin T Mansur. Do oligopolists pollute less? evidence from a restructured electricity market. *The Journal of Industrial Economics*, 55(4):661–689, 2007.

Gregory J Martin and Ali Yurukoglu. Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–2599, 2017.

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*, 2023.

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.

Nathan H Miller and Matthew C Weinberg. Understanding the price effects of the millercoors joint venture. *Econometrica*, 85(6):1763–1791, 2017.

Carlisle E Moody. Firearms and homicide. In *Handbook on the Economics of Crime*. Edward Elgar Publishing, 2010.

Terrence J Moore, Brian M Sadler, and Richard J Kozick. Maximum-likelihood estimation, the cramér–rao bound, and the method of scoring with parameter constraints. *IEEE Transactions on Signal Processing*, 56(3):895–908, 2008.

Sarah Moshary, Bradley T Shapiro, and Sara Drango. Preferences for firearms. *American Economic Review: Insights*, 7(3):340–356, 2025.

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. *American Economic Journal: Applied Economics*, 16(4):34–70, 2024.

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.

Max Pienkny, Maya Rossin-Slater, Molly Schnell, and Hannes Schwandt. The lasting impacts of school shootings on youth psychotropic drug use. *AEA Papers and Proceedings*, 114:387–93, May 2024. doi: 10.1257/pandp.20241085. URL https://www.aeaweb.org/articles?id=10.1257/pandp.20241085.

Arthur Cecil Pigou. *The Economics of Welfare*. Macmillan, 1924.

Frank P Ramsey. A contribution to the theory of taxation. *The economic journal*, 37(145): 47–61, 1927.

Mathias Reynaert. Abatement strategies and the cost of environmental regulation: Emission standards on the european car market. *The Review of Economic Studies*, 88(1):454–488, 2021.

Adam Rosenberg. Regulating firearm markets: Evidence from california. 2025.

Agnar Sandmo. Optimal taxation in the presence of externalities. *The Swedish Journal of Economics*, pages 86–98, 1975.

Terry L. Schell, Samuel Peterson, Brian G. Vegetabile, Adam Scherling, Rosanna Smart, and Andrew R. Morral. *State-Level Estimates of Household Firearm Ownership*. RAND Corporation, Santa Monica, CA, 2020. doi: 10.7249/TL354.

Bradley Shapiro, Sara Drango, and Sarah Moshary. Firearm policy buyers would accept. *Available at SSRN 6514240*, 2026.

Rosanna Smart. Firearm and ammunition taxes. Technical report, 2021. URL https://www.rand.org/research/gun-policy/analysis/essays/firearm-and-ammunition-taxes.html.

Rosanna Smart, Andrew R. Morral, James P. Murphy, Rupa Jose, Amanda Charbonneau, and Sierra Smucker. *The Science of Gun Policy: A Critical Synthesis of Research Evidence on the Effects of Gun Policies in the United States, Fourth Edition*. RAND Corporation, Santa Monica, CA, 2024. doi: 10.7249/RRA243-9.

Sierra Smucker, Max Griswold, Amanda Charbonneau, Rose Kerber, Terry L Schell, and Andrew R Morral. *Using National Instant Criminal Background Check Data for Gun Policy Analysis: A Discussion of Available Data and Their Limitations*. RAND, 2022.

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.

Marcella Alsan Joshua Schwartzstein Stefanie Stantcheva. The universal pursuit of safety and the demand for (lethal, non-lethal or no) guns. 2025.

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.

U.S. Bureau of Economic Analysis. Input-output accounts data for ammunition, arms, ordnance, and accessories manufacturing. 2017.

Frank Verboven. International price discrimination in the european car market. *The RAND Journal of Economics*, pages 240–268, 1996.

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

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

	(1) Adult Pop.	(2) Potential Gun Buyers	(3) Gun Buyers	(4) Handgun Buyers	(5) Longgun Buyers
Female	0.521	0.430	0.085	0.107	0.048
Fraction White in Zipcode	0.718	0.738	0.813	0.804	0.828
Poverty Rate in Zipcode	0.106	0.100	0.081	0.083	0.078
Conservative Vote Share (2016)	0.365	0.380	0.440	0.437	0.444
Density of Zipcode	1,940	1,677	666	712	587
Median Zipcode Income	87,263	87,715	88,528	88,031	89,381
Fraction BA+ in Zipcode	0.438	0.433	0.402	0.399	0.407

Table 1: Demographics of Gun Buyers in Massachusetts

Table reports the average demographic characteristics across different sub-populations of adults in Massachusetts. Rows are demographic characteristics. Columns are sub-populations. Each value is a weighted average of a demographic characteristic across the population of gender-by-zip code cells. Columns weight by (1) adult population; (2) consumer market size, measured by the number of adults who own or could see themselves owning a gun; (3) the observed number of firearm purchases in our sample; (4) the observed number of handgun purchases in our sample; and (5) the observed number of long gun purchases in our sample. Across rows, the demographic for each gender-by-zip code cell is measured, respectively, by a binary indicator for female gender, the share of zip code population that is white, the share of zip code population below the federal poverty line, the vote share of all conservative candidates in the 2016 presidential election, population per square mile, median zip code income, the share of zip code population that has a bachelor's degree.

Weapon Class:	Handguns		Long guns	
Aggregation:	Model-Years	Transactions	Model-Years	Transactions
	(1)	(2)	(3)	(4)
Caliber (Inches)	0.357	0.356	0.444	0.424
	(0.084)	(0.056)	(0.217)	(0.219)
Barrel Length (inches)	4.294	3.735	22.447	20.315
	(1.779)	(0.962)	(4.385)	(3.847)
High Capacity Weapon	0.185	0.329	0.152	0.283
	(0.388)	(0.470)	(0.359)	(0.450)
Shotgun	0.000	0.000	0.379	0.325
	(0.000)	(0.000)	(0.485)	(0.468)
In Active Production (New)	0.356	0.742	0.391	0.617
	(0.479)	(0.438)	(0.488)	(0.486)
MSRP (2022$)	1034.91	734.55	1657.83	1040.95
	(751.14)	(439.02)	(2816.27)	(1159.78)
Number of Gun Models	2,835		4,781	
Number of Transactions/Year	71,254		33,450	

Table 2: Firearm Characteristics

Table reports the average value and (standard deviation) of firearm characteristics, across firearm model-years in the Blue Book of Gun Values that we match to at least one transaction in Massachusetts during our sample. Rows are characteristics of a firearm model-year. Columns are different aggregations across firearm model-years, separately for handguns in Columns (1)–(2) and long guns in Columns (3)–(4). Each value is a weighted average (or standard deviation) of a characteristic across firearm model-years. Columns weight by (1) unity, for handguns; (2) the number of purchases of each handgun model-year from firearm dealers in Massachusetts during our sample; and (3, 4) analogous for long guns. Across rows, the characteristic for each firearm model is measured, respectively, by (1) the caliber of ammunition fired by the firearm model-year, in inches; (2) the length of the firearm model-year's barrel, in inches; (3) a binary indicator for a firearm model-year that has a high-capacity magazine; (4) a binary indicator for a firearm model-year that is a shotgun; (5) a binary indicator for a firearm model-year that is under active production by manufacturers; (6) the firearm model-year's MSRP in 2022 dollars, measured only for firearms under active production. Number of gun models is the number of distinct firearm models used in computing the weighted statistics across firearm model-years. Number of transactions/year is the average number of yearly firearm transactions used in computing the weighted statistics across firearm model-years.

	(1) Log(Purchases)	(2) Log(Purchases)	(3) Log(Purchases)
Log(MSRP)	-0.834***	-1.557*	-2.136**
	(0.317)	(0.860)	(0.831)
Observations	2436	2436	2436
Gun Model FE	Yes	Yes	Yes
Year FE	Year	Year	Class x Year
# Potential Instruments		2772	2772
# Selected Instruments		6	7
Sup-Score Test Statistic		12.44	12.57
First Stage Sig. (p-value)		0.0000	0.0000
Method	OLS	IVLASSO	IVLASSO

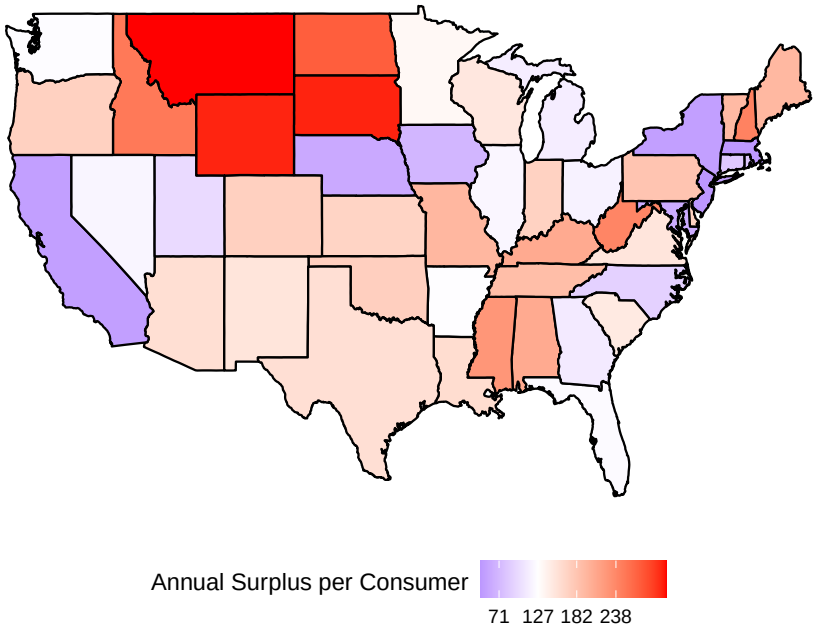
Table 3: Estimated Price Elasticity of Demand for Firearms

Table displays the estimated own-price elasticity of demand from the IVLASSO routine described in Section 4.1 applied to a sample of 2,436 firearm model-years. The sample is restricted to firearm models with at least 100 purchases in Massachusetts between 2016–2022. Dependent variable is the log of a firearm model's purchase quantity in Massachusetts each year. Regressor is the log of a firearm's MSRP each year. Column (1) reports estimates from an OLS regression with fixed effects for gun model j and year t. Columns (2) and (3) report estimates that instrument for MSRP using our cost shock from Section 4.1, based on manufacturer-specific exposure to metal commodity prices. Column (2) includes fixed effects for gun model j and year t. Column (3) includes fixed effects for gun model j and class-year $c(j), t$. In Columns (2) and (3), #Potential Instruments is the number of manufacturer-commodity pairs assumed to be valid cost-shock instruments, and #Selected Instruments is the number of these manufacturer-commodity pairs chosen by the IVLASSO procedure to optimize estimation performance. Sup-Score Test Statistic is the statistic from a joint test against the irrelevance of all the instruments selected by the IVLASSO procedure, with the associated p-value below (i.e., low p rejects the null of irrelevant instruments, adjusting for the LASSO selection procedure). Standard errors robust to heteroskedasticity in parentheses. * denotes $p < .1$, ** denotes $p < .05$, *** denotes $p < .01$.

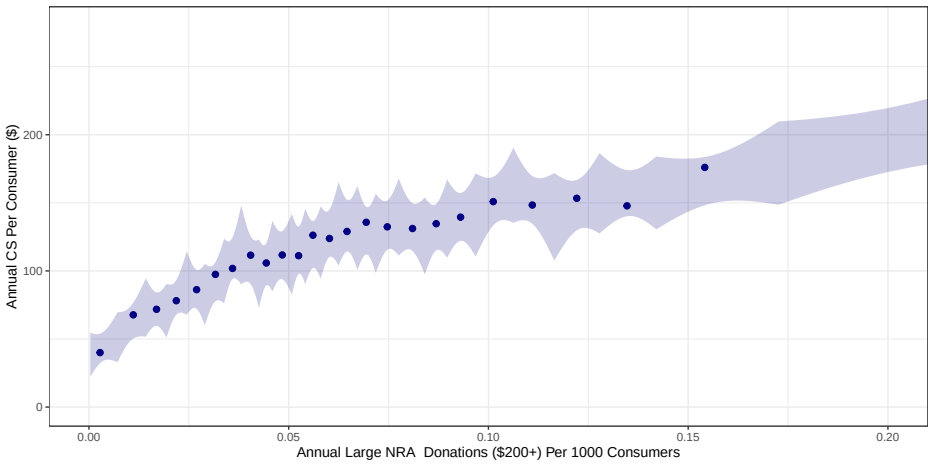
Object	Value	Source
Panel A: Gun Stock law of Motion		
$\text{Households}_{s,2015}$	———	2015 ACS, 5-year estimate
$Pr(\text{Adult has gun in HH})_{s,2015}$	———	Schell et al. (2020)
$\text{Gun owners per HH with gun}_{2015}$	1.705	Azrael et al. (2017)
$\text{Guns per Gun Owner}_{c,s,2015}$	———	Azrael et al. (2017)
$Q_{c,s,2015}$: Baseline gun stock	———	
φ: Scrap rate	0.015 (0.004)	Cook and Ludwig (1996), Azrael et al. (2017), ATF firearms commerce report
Panel B: Distribution of Gun Homicides		
κ: Elasticity of gun homicide wrt HH gun ownership	0.294 (0.078)	Duggan (2001), Cook and Ludwig (2006)
$Pr(\text{Shot handgun} \mid \text{Gun homicide})$	0.899	Morgue records from NVDRS and NCHS
χ_h/χ_l: Homicide "efficiency" of handguns relative to long guns	8.902 (0.424)	
$\chi_l\zeta_{s,t}$: State-year homicide shock	———	
Panel C: Average marginal homicides from firearm purchase		
Handgun	3.347×10^{-5}	
Long gun	3.609×10^{-6}	
Share handguns among purchases	0.761	
Average marginal homicide	2.634×10^{-5}	
Panel D: Normative conversion		
Welfare cost, 1 gun homicide	\$9.5m	Peterson et al. (2021)
Yearly discount factor	0.95	

Table 4: Calibration sources and parameter estimates for homicide model

Table displays data source, calibration targets, and parameter values from the model of firearm homicides in Section 6. For heterogeneous terms "—" suppresses state-specific values in Panel A and state-year specific values in Panel B. Standard errors in parentheses are computed from the gradient of the calibration objective with respect to the parameter estimates, separately for the scrap rate φ in Panel A and the homicide parameters $\kappa, \chi_h/\chi_l$ in Panel B. Details in Appendix E.



(a) Consumer Surplus Across the U.S.



(b) Consumer Surplus Correlates with NRA Donations

Figure 1: Consumer Surplus from the Firearms market in the U.S.

Figure displays consumer surplus from the U.S. firearms industry. Panel (a) displays the annual dollar value of consumer surplus (CS) per consumer in each state, averaged over 2016-2022. White represents the population-weighted mean CS per consumer-year in the U.S. Scale ticks represent population-weighted standard deviations of CS per consumer-year. Panel (b) displays a binned scatterplot from a dataset of zip codes, showing average annual CS per consumer (vertical axis) against the average annual number of donations $\geq$ \$200 to the NRA, per 1,000 consumers. The procedure of Cattaneo et al. (2019) is used to select the number and location of points and to form 95% confidence intervals.

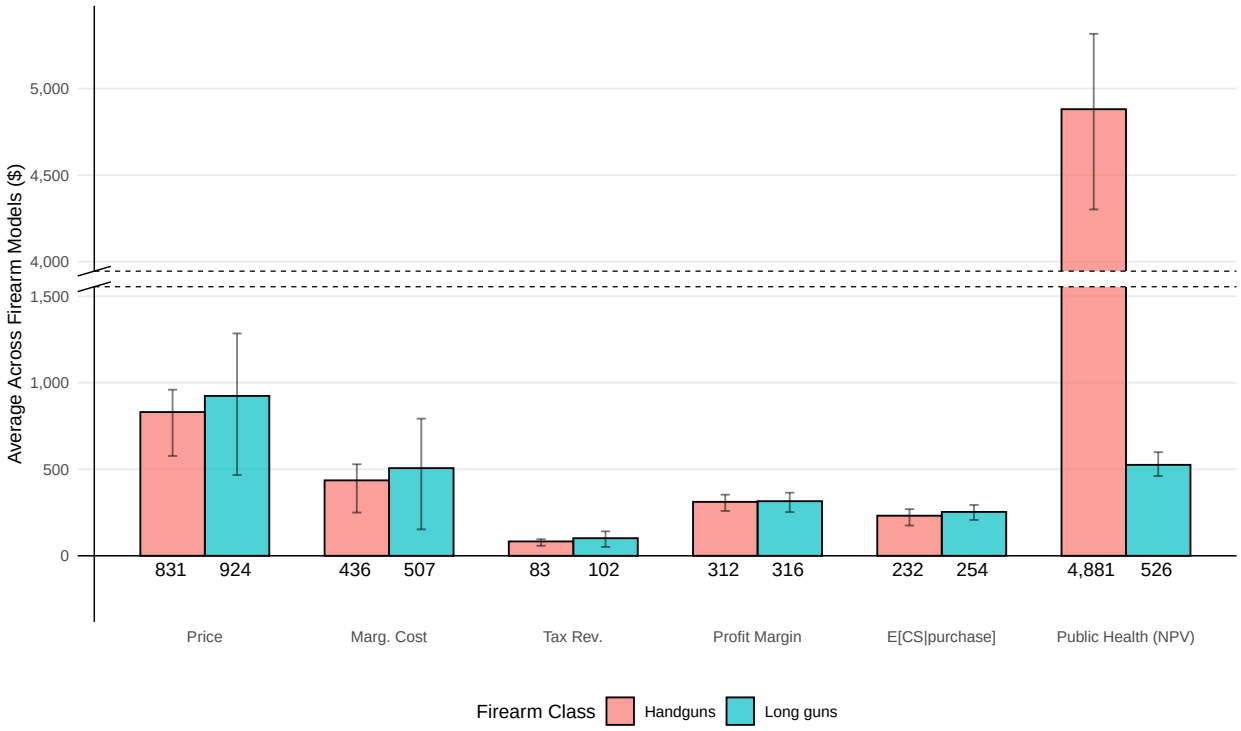
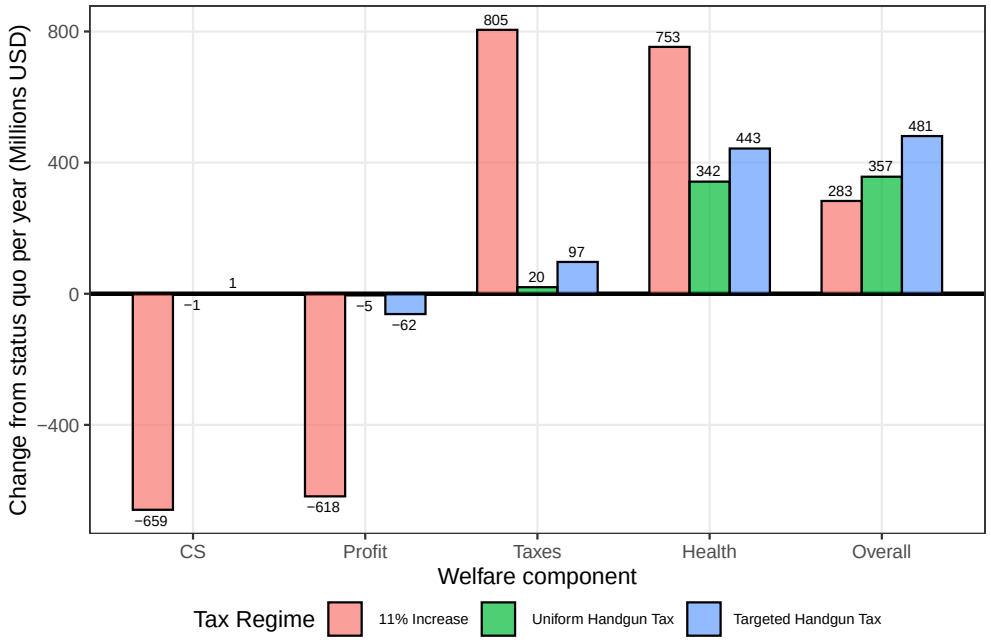
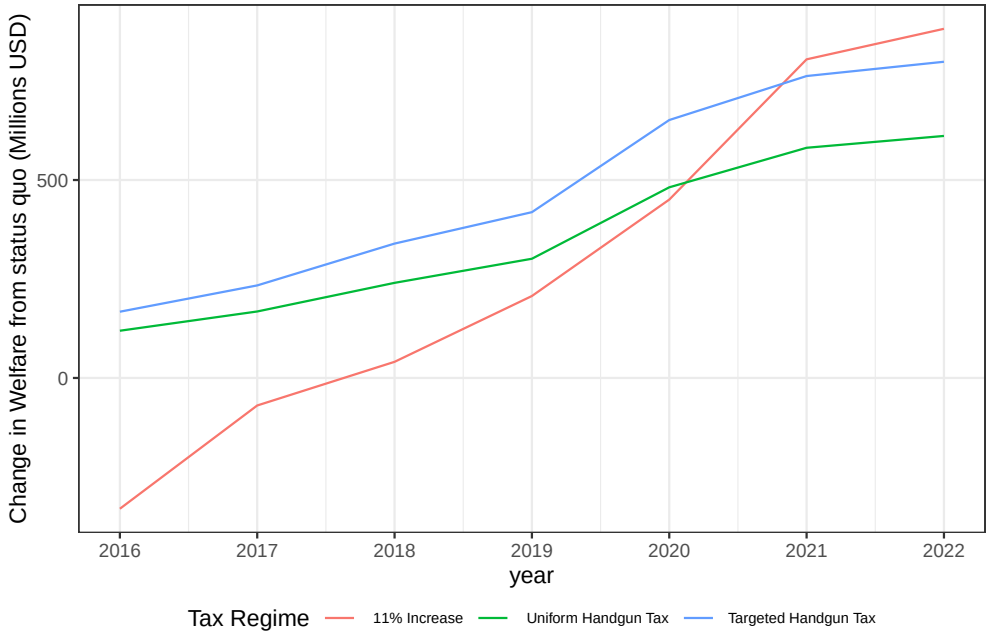


Figure 2: Distribution of Surplus from Firearm Purchase

Figure displays distribution of surplus in dollars from the average marginal firearm purchase in each class. Bars are the unweighted-average of a statistic across firearm-years within a class. Bands are the unweighted 25th and 75th percentiles of the statistics. Price is the firearm's MSRP $p_{j,t}$. Tax Revenue is the revenue to the federal government generated by the firearm's sale under status quo regulation $v_j p_{j,t}$. Marginal cost is our estimate of the marginal production cost $c_{j,t}$. Profit margin is our estimate of price, less taxes and marginal cost $\mu_{j,t} = p_{j,t}(1 - v_j) - c_{j,t}$. By construction, the sum of tax revenue, marginal cost, and profit margin equals the price. Consumer surplus measures the average value to consumers from purchasing the firearm, conditional on that firearm being the best alternative in their choice set. Public health is equal to the net present value of expected lifetime gun homicides caused by a firearm purchase, converted to dollars via a value of a statistical life. Vertical axis suppresses \$1,500–\$4,000 for readability.



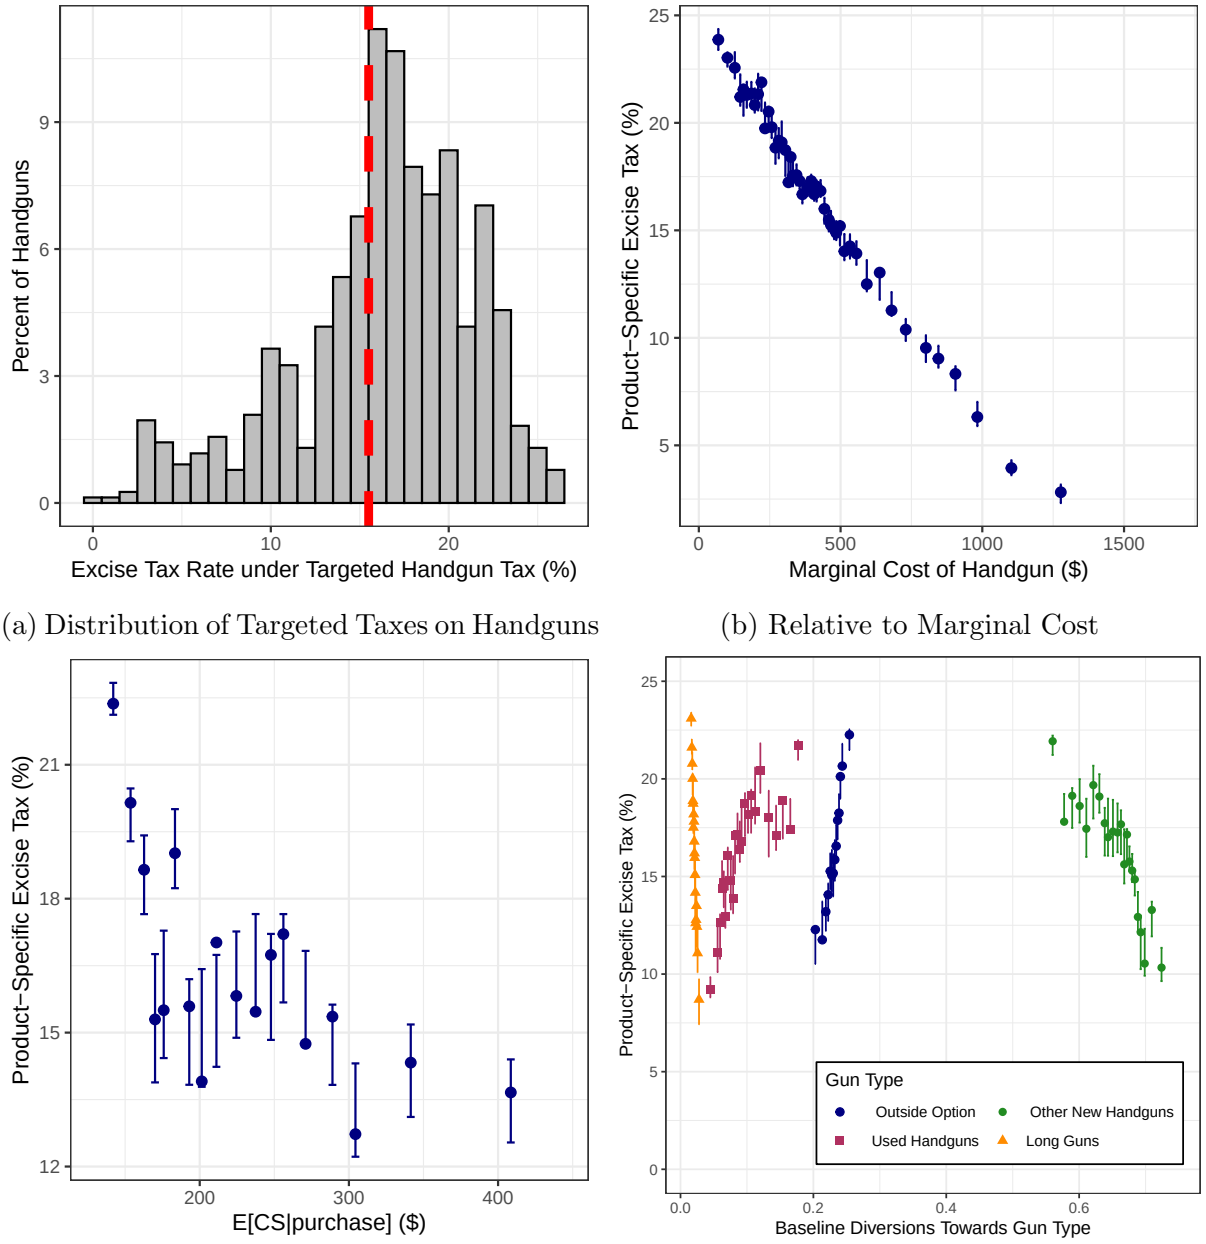
(a) Average Welfare Effects



(b) Overall Welfare Effects By Year

Figure 3: Equilibrium Welfare Effects of Firearms Tax Policies

Figure shows the change in dollars of welfare nationwide, comparing the status quo to the counterfactual firearm taxes in Section 7. Counterfactuals tax systems: (orange) homogeneously raise the tax on all firearms by 11 percentage points, (green) set the tax on all long guns to 0% and on all handguns to 15.5%, (blue) set the tax on all long guns to 0% and a set heterogeneous tax on each handgun-year by approximating the constrained-optimal tax in Equation (3). Panel (a), shows the average in-sample annual change in each of the four welfare components from Equation (1) and the change in overall welfare (their sum). Panel (b), shows the annual change in overall welfare.



(a) Distribution of Targeted Taxes on Handguns

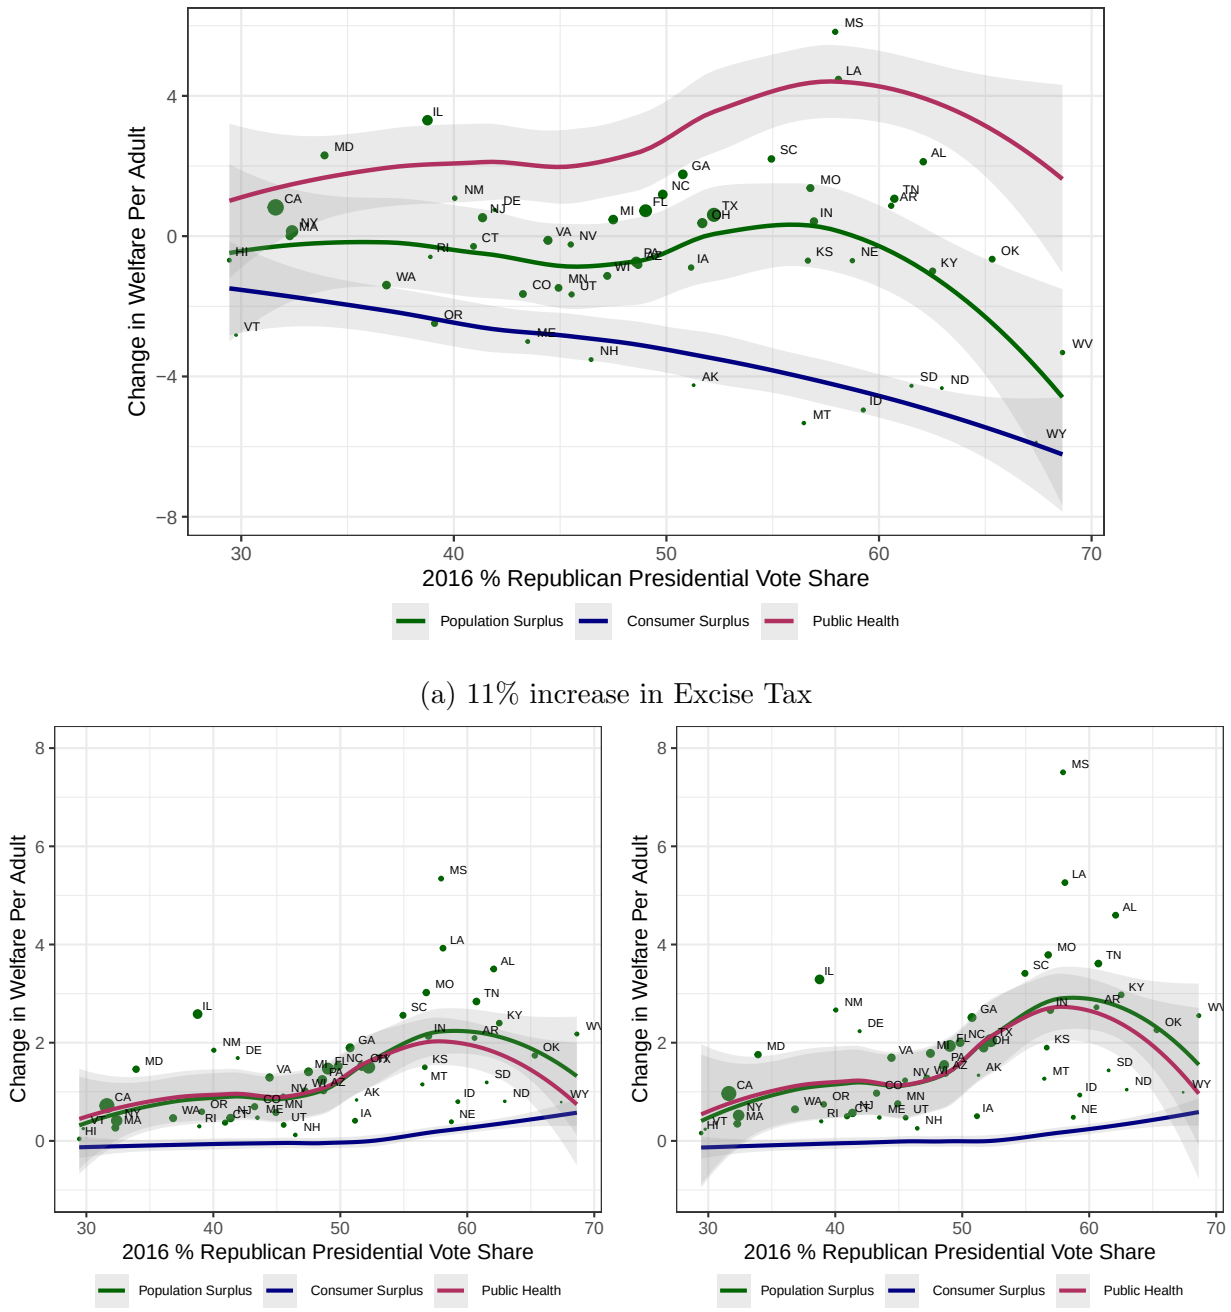
(b) Relative to Marginal Cost

(c) Relative to Consumer Surplus

(d) Relative to Substitution Patterns

Figure 4: Incidence of the Targeted Handgun Tax Across Handgun Models

Panel (a) displays the distribution of excise tax rates across handgun-years under the targeted handgun tax, relative to the uniform handgun tax of 15.5% (dashed red line). Panel (b) displays the targeted handgun tax rates as a binned scatterplot across handgun-years, relative to the estimated marginal cost of production of each handgun-year. Panel (c) displays the targeted handgun tax rates as a binned scatterplot across handgun-years, relative to the expected consumer surplus of buyers of each handgun-year under the status quo equilibrium. Panel (d) displays the targeted handgun tax rates as a binned scatterplot across handgun-years, relative to each handgun-year's estimated diversion ratio towards other types of options in the choice set. Each handgun-year contributes once, independently to each series in Panel (d). In Panels (b)–(d), the procedure of Cattaneo et al. (2019) is used to select the number and location of points and to form 95% confidence intervals, separately for each series.



(a) 11% increase in Excise Tax

(b) Uniform Handgun Tax

(c) Targeted Handgun Tax

Figure 5: Average Welfare Effects by Political Preferences

Figure shows changes in dollars of average annual surplus per-capita relative to the Republican vote share in the 2016 Presidential Election, for each U.S. state. Each panel compares the status quo to a counterfactual firearm tax from Section 7: (a) homogeneously raise the tax on all firearms by 11 percentage points, (b) set the tax on all long guns to 0% and on all handguns to 15.5%, (c) set the tax on all long guns to 0% and a set heterogeneous tax on each handgun-year by approximating the constrained-optimal tax in Equation (3). Population Surplus is defined as the sum of consumer surplus and the net present dollar value of improvements in public health. Horizontal is the state-level share of votes for Donald Trump in the 2016 U.S. Presidential Election. Vertical axis is the change in the welfare component from the counterfactual tax. For each welfare component, we plot a LOESS fit of the data, weighted by state population. Points correspond to each state's estimated change in population surplus, with size proportional to adult population.

A Optimal Excise Tax Derivation

This section derives the optimal specific excise tax formula presented in Section 2. It also derives the optimal *ad valorem* excise tax formula and considers the structure of optimal taxes under perfect competition. Finally, it generalizes the optimal tax formula to a planner whose marginal value of public funds differs from one.

Proof of Proposition 1.

Set up In this proof, we suppress the market index t and interpret j as an index across distinct product-markets. The transaction quantity of j is q_j, with all quantities collected in $\vec{q}$.

Let $\Phi(\vec{q})$ denote the negative externality from consumption. The externality is a function of all purchase quantities $\vec{q}$. It has partial derivatives $\phi_j \equiv \partial\Phi/\partial q_j$, where $\phi_j \in \mathbb{R}$.

Let τ_j denote the specific tax on j. With specific taxes, government revenue is $\mathcal{G} = \sum_j \tau_j q_j$.

The firm (potentially multi-product and/or multi-market) with ownership of j earns profit from its sale $\pi_j \equiv q_j(p_j - \tau_j) - C_j(q_j)$. Production cost $C_j(q_j)$ may be non-linear, but has strictly positive marginal costs $c_j = \partial C_j/\partial q_j > 0$ that are separable across the index $\partial C_j/\partial q_j \partial q_k = 0$. Recall that the firm's post-tax profit margin on j is $\mu_j \equiv p_j - \tau_j - c_j$.

Each consumer operates in exactly one market, exogenously determined, and enjoys surplus from the purchase of their most-preferred item in their market. Since consumers optimize, their change in aggregate surplus CS due to a marginal price change is $\partial CS/\partial p_j = -q_j$.

Planner We characterize the planner's solution by analyzing how marginal changes to the tax system are *passed through* to equilibrium prices $P_{j,k} \equiv \partial p_j/\partial \tau_k$ and quantities $S_{j,k} \equiv \partial q_j/\partial \tau_k = \sum_l \partial q_j/\partial p_l \cdot \partial p_l/\partial \tau_k = \sum_l D_{j,l} \cdot P_{l,k}$.

For instance, a change in the tax on product k affects consumer surplus by

$$\frac{\partial}{\partial \tau_k} CS = \sum_j \frac{\partial CS}{\partial p_j} \frac{\partial p_j}{\partial \tau_k} = \sum_j -q_j \cdot P_{j,k}.$$

Similarly, a change in the tax on product k affects surplus attributable to firms and the

government (profits and tax revenue, respectively) by

$$\begin{aligned}\frac{\partial}{\partial \tau_k}(\Pi + \mathcal{G}) &= \frac{\partial}{\partial \tau_k}\left(\sum_j p_j q_j - C_j(q_j) - \tau_j q_j + \tau_j q_j\right) \\ &= \sum_j p_j \frac{\partial q_j}{\partial \tau_k} + q_j \frac{\partial p_j}{\partial \tau_k} - c_j \frac{\partial q_j}{\partial \tau_k} \\ &= \sum_j \underbrace{(p_j - c_j)}_{\equiv \tau_j + \mu_j} \cdot S_{j,k} + \underbrace{\sum_j q_j \cdot P_{j,k}}_{\equiv -\partial CS/\partial \tau_k}.\end{aligned}$$

This condition reflects two important consequences of the planner's objective in Equation (1). In the first line, tax revenue $\tau_j q_j$ is a pure transfer of welfare from firms to the government. In the last line, revenue from higher prices $q_j P_{j,k}$ is a transfer from consumers to other market participants.

Formally, the planner solves

$$\begin{aligned}\max_{\vec{\tau}} \quad & CS(\vec{p}) + \sum_j \Pi_j(\vec{p}, \vec{q}) + \sum_j \mathcal{G}(\vec{p}, \vec{q}) - \Phi(\vec{q}) \\ \text{subject to:} \quad & CS(\vec{p}) \geq CS_0,\end{aligned}$$

where both prices $\vec{p}(\vec{\tau})$ and quantities $\vec{q}(\vec{\tau})$ are functions of the tax schedule $\vec{\tau}$. The associated lagrangian places a multiplier λ on the consumer surplus constraint:

$$\max_{\vec{\tau}} \quad CS(\vec{p}) + \sum_j \Pi_j(\vec{p}, \vec{q}) + \sum_j \mathcal{G}(\vec{p}, \vec{q}) - \Phi(\vec{q}) + \lambda(CS(\vec{p}) - CS_0)$$

Taking the first order condition with respect to τ_k, and substituting the derivatives above, yields

$$\begin{aligned}0 &= \sum_j \mu_j S_{j,k} - \sum_j \phi_j S_{j,k} - \lambda \sum_j q_j P_{j,k} \\ \implies \lambda \sum_j q_j P_{j,k} &= \sum_j (\tau_j + \mu_j - \phi_j) S_{j,k}\end{aligned}$$

We express the first order conditions of the lagrangian in matrix form. Let D be the matrix of demand price derivatives (e.g. $D_{j,k} = \partial q_j / \partial p_k$), with P and S the conformable matrices of tax passthrough to price and quantity, such that $S = D \cdot P$. Stacking the first

order conditions across τ_k gives

$$\begin{aligned} \lambda P' \cdot \vec{q} &= S' \cdot (\vec{\tau} + \vec{\mu} - \vec{\phi}) \\ \implies \lambda D'^{-1} \cdot \vec{q} &= \vec{\tau} + \vec{\mu} - \vec{\phi}, \end{aligned} \tag{A.1}$$

where the second line premultiplies both sides by $S'^{-1} = D'^{-1}P'^{-1}$.

Monopoly Consider the problem now of a monopolist, who wishes to maximize (static) profits. Their optimization problem is:

$$\max_{\vec{p}_f} \sum_j (p_j - \tau_j) q_j - C_j(q_j).$$

The first order condition for any element p_k of $\vec{p}_f$ is

$$0 = q_k + \sum_j (\mu_j - \tau_j) \frac{\partial q_j}{\partial p_k}$$

Stacking these first order conditions across all products yields the familiar first order condition in matrix form

$$\begin{aligned} \vec{0} &= \vec{q} + (D)'(\vec{\mu} - \vec{\tau}) \\ \implies \vec{\mu}^M &= -(D)'^{-1}\vec{q}, \end{aligned}$$

This monopolist markup $\vec{\mu}^M$ corresponds exactly to the term on the left-hand side of Equation A.1. We note that because the monopolist faces no competition, the particular conduct assumption (e.g. Cournot versus Bertrand pricing) on the monopolist does not affect the solution, so long as it is a static profit maximizer. Rearranging yields the optimal tax

$$\tau^* = \vec{\phi} - \vec{\mu} - \lambda \vec{\mu}^M$$

Or, for each element/product j

$$\tau_j^* = \phi_j - \mu_j - \lambda \mu_j^M.$$

Which concludes the proof. □

In our empirical setting, ad valorem excise taxes are used instead of specific taxes. We

cover this case in the following proposition.

Proposition 2. *Suppose the social planner optimizes welfare as defined in Equation (1), choosing ad valorem excise taxes $v_{j,t}$, and is subject to a consumer surplus constraint as in Equation (2). Then the optimal ad valorem excise tax on product j in market t is:*

$$v^*_{j,t} = \frac{\phi_{j,t} - \mu_{j,t} - \lambda \cdot \mu^M_{0,j,t}}{p_{j,t}} \tag{A.2}$$

Where $\phi_{j,t}$ is the magnitude of the (negative) marginal externality from consuming the product, $\mu_{j,t}$ is the margin charged by suppliers on their products, $\mu^M_{0,j,t}$ is the margin that would be charged under monopoly ownership with no taxes $v_{j,t} = 0$, $p_{j,t}$ is the price of product j in market t, and λ is the shadow cost (in welfare units) of the consumer surplus constraint.

Proof. The characterization of the planner's problem under ad valorem taxes is analogous to the case of specific taxes. The key difference between these systems is the structure of the tax-inclusive markup

$$p - c_j = \tau_j + \mu_j = p_j \cdot v_j + \mu_j.$$

Making this substitution into Equation (A.1) yields the planner's first-order conditions in matrix form:

$$\lambda D'^{-1} \cdot \vec{q} = \vec{p} \odot \vec{v} + \vec{\mu} - \vec{\phi}. \tag{A.3}$$

where $\odot$ denotes the Hadamard product.

The monopolist's problem is parallel to the problem presented in the above proof. The first order condition is:

$$\begin{aligned} 0 &= (1 - v_k)q_k + \sum_j ((1 - v_j)p_j - c_j)\frac{\partial q_j}{\partial p_k} \\ \implies -(\vec{1} - \vec{v}) \odot \vec{q} &= (D')(\vec{p} \odot (\vec{1} - \vec{v}) - \vec{c}) \\ \implies \vec{\mu}^M &= -(D')^{-1}((\vec{1} - \vec{v}) \odot \vec{q}) \end{aligned}$$

Because the Hadamard product does not commute (unless D is diagonal), the monopoly markup under ad valorem taxation does not appear in Equation (A.3). Instead, define $\vec{\mu}^M_0$

as the monopoly markup under zero ad valorem taxes ($\vec{v} = \vec{0}$), whence

$$\vec{\mu}_0^M = -D'^{-1}(\vec{q}),$$

and this is identical in form to the monopoly markup under any specific tax system $\vec{\tau}$.

Substituting the monopoly markup under zero ad valorem taxes $\vec{\mu}_0^M$ into the planner's first order condition in Equation (A.3) and rearranging yields

$$p \odot \vec{v}^* = \vec{\phi} - \vec{\mu} - \lambda \vec{\mu}_0^M$$

Or, for each element/product j:

$$v_j^* = \frac{\phi_j - \mu_j - \lambda \mu_{0,j}^M}{p_j}.$$

Which completes the proof. □

Finally, we also implement the optimal tax in our assumption of competitive pricing, where $p_{j,t}(1 - v_{j,t}) = c_{j,t}$. That is, profits are zero, and firms produce at cost. A corollary to Proposition 1 establishes its form. All we need to edit is $\mu_j = 0$ since profits are by construction zero in this conduct model. The optimal tax under ad valorem taxes is once again very similar, only requiring division by $p_{j,t}$ as a modification.

Corollary 1. *Suppose that firms provide firearms at cost, and the social planner optimizes welfare as the sum of consumer surplus CS, government revenue $\mathcal{G}$, and negative externalities associated with consumption Φ, choosing specific excise taxes τ_j, and is subject to a consumer surplus constraint as in Equation 2. Then the optimal specific excise tax on product j is:*

$$\tau_j^* = \phi_j - \lambda \cdot \mu_j^M \tag{A.4}$$

Where ϕ_j is the magnitude of the (negative) marginal externality from consuming the product, μ_j^M is the margin that would be charged under a monopoly or cartel conduct, and λ is the shadow price (in welfare units) of the consumer surplus constraint.

Proof. To establish the result, simply set $\vec{\mu} = 0$, which occurs by construction. The monopoly margin is then interpreted as the markups assuming they faced the same demand curve as the (competitive) firms. □

Throughout the proofs, we use Equation (1) to measure welfare, which values a dollar of government revenue at par with a dollar of private surplus, implicitly assuming that the

marginal value of public funds (MVPF) equals one. A significant literature has established that government revenue may be socially more or less beneficial depending on its use case (Hendren and Sprung-Keyser 2020). The final proposition of this section generalizes Proposition 1 to a planner who values each dollar of government revenue at η dollars of welfare.

Proposition 3. *Suppose the social planner chooses specific excise taxes $\tau_{j,t}$ to solve the following optimization problem:*

$$\begin{aligned} &\max_{\tau_{j,t}} \quad CS + \Pi + \eta \cdot \mathcal{G} - \Phi \\ &s.t.\ CS(\vec{\tau}) \geq CS_0, \end{aligned} \tag{A.5}$$

where $\eta > 0$ is the marginal value of public funds. Then the constrained-optimal specific tax on product j in market t satisfies:

$$\tau^*_{j,t} = \frac{1}{\eta}\Big(\phi_{j,t} - \mu_{j,t} - \lambda \cdot \mu^M_{j,t}\Big) + \frac{\eta - 1}{\eta} \cdot \tau^L_{j,t} \tag{A.6}$$

Where $\phi_{j,t}$, $\mu_{j,t}$, and $\mu^M_{j,t}$ are the marginal externality, profit margin, and monopoly margin defined in Proposition 1, λ is the shadow cost of the consumer surplus constraint under the welfare criterion in Equation (A.5), and $\tau^L_{j,t}$ is the tax that a government seeking solely to maximize tax revenue (the so-called "Leviathan tax") would set on product j in market t.

Proof. We adopt the setup and notation from the proof of Proposition 1, again suppressing the market index t.

When $\eta \neq 1$, tax revenue is no longer a pure transfer between firms and the government, so profits and revenue enter the planner's first order condition separately. Their derivatives with respect to the tax on product k are

$$\begin{aligned} \frac{\partial \Pi}{\partial \tau_k} &= \frac{\partial}{\partial \tau_k}\left(\sum_j q_j(p_j - \tau_j) - C_j(q_j)\right) = \sum_j \left(\mu_j \cdot S_{j,k} + q_j \cdot P_{j,k}\right) - q_k, \\ \frac{\partial \mathcal{G}}{\partial \tau_k} &= \frac{\partial}{\partial \tau_k} \sum_j \tau_j q_j = q_k + \sum_j \tau_j \cdot S_{j,k}. \end{aligned}$$

As before, the lagrangian places a multiplier λ on the consumer surplus constraint. Taking

its first order condition with respect to τ_k, and substituting the derivatives above, yields

$$\begin{aligned} 0 &= -(1+\lambda)\sum_j q_j P_{j,k} + \sum_j (\mu_j S_{j,k} + q_j P_{j,k}) - q_k + \eta\Big(q_k + \sum_j \tau_j S_{j,k}\Big) - \sum_j \phi_j S_{j,k} \\ &= -\lambda \sum_j q_j P_{j,k} + (\eta - 1)\, q_k + \sum_j (\eta\tau_j + \mu_j - \phi_j)\, S_{j,k}. \end{aligned}$$

Stacking the first order conditions across τ_k gives, in matrix form,

$$\lambda P' \cdot \vec{q} - (\eta - 1)\, \vec{q} = S' \cdot (\eta\vec{\tau} + \vec{\mu} - \vec{\phi}). \tag{A.7}$$

Relative to the first order conditions in the proof of Proposition 1, the excess value of public funds $\eta - 1$ introduces the mechanical revenue effect of taxation $\vec{q}$, no longer offset by an equal-sized loss in profits. We interpret this first order condition via an alternative revenue-maximizing problem of the government.

Leviathan Consider the problem of a government that seeks solely to maximize its tax revenue—the "Leviathan" government of Brennan and Buchanan (1980):

$$\max_{\vec{\tau}} \sum_j \tau_j q_j.$$

The first order condition for any tax τ_k is

$$0 = q_k + \sum_j \tau_j \cdot S_{j,k},$$

and stacking these first order conditions across all products yields the revenue-maximizing tax system in matrix form

$$\vec{0} = \vec{q} + S' \cdot \vec{\tau} \implies \vec{\tau}^L = -(S')^{-1} \cdot \vec{q}.$$

Premultiplying Equation (A.7) by $S'^{-1} = D'^{-1}P'^{-1}$, then substituting the monopoly markup $\vec{\mu}^M = -D'^{-1} \cdot \vec{q}$ and the Leviathan tax $\vec{\tau}^L = -(S')^{-1} \cdot \vec{q}$, yields

$$\begin{aligned} \eta\vec{\tau} + \vec{\mu} - \vec{\phi} &= \lambda D'^{-1} \cdot \vec{q} - (\eta - 1)(S')^{-1} \cdot \vec{q} \\ &= -\lambda\vec{\mu}^M + (\eta - 1)\, \vec{\tau}^L. \end{aligned}$$

Rearranging yields the optimal tax

$$\vec{\tau}^* = \frac{1}{\eta}\left(\vec{\phi} - \vec{\mu} - \lambda\vec{\mu}^M\right) + \frac{\eta - 1}{\eta} \cdot \vec{\tau}^L$$

Or, for each element/product j

$$\tau_j^* = \frac{1}{\eta}\left(\phi_j - \mu_j - \lambda\mu_j^M\right) + \frac{\eta - 1}{\eta} \cdot \tau_j^L.$$

Which concludes the proof. □

The optimal tax in Proposition 3 is a weighted average of two tax systems: the second-best tax of Proposition 1 and the revenue-maximizing Leviathan tax, with weight $\frac{\eta-1}{\eta}$ on the latter. When $\eta = 1$, the formula collapses to the second-best tax of Proposition 1. When $\eta > 1$, each dollar of revenue purchases more than a dollar of welfare, pulling the optimal tax system away from the second-best schedule and towards the Leviathan schedule, which it approaches as $\eta \to \infty$. The Leviathan tax itself has a familiar structure: under complete passthrough of taxes to prices ($P = I$), $\vec{\tau}^L = \vec{\mu}^M$, so the revenue-maximizing government sets taxes that replicate the margins of a monopolist over the entire industry.

B Details on Setting and Data

B.1 Additional Datasets

We describe in this section the sources and data collection methods of supplementary datasets not described in the main text.

FEC Contributions Data We download all recorded individual contributions from the Federal Election Commission[55] to three Political Action committees associated with the NRA: The National Rifle Association of America Political Victory Fund, NRA Victory Fund, Inc., and the National Rifle Association Institute for Legislative Action. We collapse these to the ZCTA-by-donation year level to compare the incidence of consumer surplus and counterfactual taxes across zip codes, according to political activity.

[55] Link: https://www.fec.gov/data/browse-data/

BLS Commodity Prices We download the price indices of each metal commodity used for our cost shock price instruments directly from the BLS website.[56] We collapse each monthly time series of a commodity's price index to its average value by year, then take a log to enter the IV-LASSO regression.

CPI To deflate prices to 2022 $ throughout our analysis, we use the CPI for the entire United States.[57]

National Firearms Survey The 2015 and 2019 National Firearms Surveys, which form the basis of our local gun stock measurements and auxiliary survey evidence on gun owner patterns, were shared with us by authors of the surveys. In both surveys, we apply the provided survey weights for computation of any statistics.

POC Status and Physical Transfers We hand-collected the Point of Contact (POC) and Brady exempt status of each state by year based on the annually filed FBI NICS operations reports from 1998 to 2024, which listed the exact month each state switched their reporting status.[58] In addition, we downloaded (where available) data from states on ground-truth physical firearms transfers or closer proxies to sales (e.g. actual non-duplicated background checks in POC states), which we use in place of adjusted NICS checks. These are:

1. California: Annual physical transfers are collected from California's Dealer Record of Sale Transaction database.[59]

2. Colorado: Annual Data on Firearm Background Checks performed by the Colorado Department of Public Health.[60]

3. Hawaii: all physical transfers of Firearms in Hawaii were documented by an annual report prepared by the Attorney General.[61]

4. Pennsylvania: Annual data on physical transfers is acquired from Johnson et al. (2023) (See Footnote 11).

[56] Data is downloaded from https://download.bls.gov/pub/time.series/wp/

[57] Data downloaded from FRED: https://fred.stlouisfed.org/series/CPALTT01USA661S

[58] Downloaded from here: https://www.fbi.gov/how-we-can-help-you/more-fbi-services-and-information/library

[59] Source: https://oag.ca.gov/sites/all/files/agweb/pdfs/firearms/forms/dros-chart-2018.pdf

[60] Source: https://cdphe.colorado.gov/resource-bank-for-firearm-violence-prevention-in-colorado

[61] An example report is located here: https://ag.hawaii.gov/cpja/files/2026/04/Firearm-Registrations-in-Hawaii-2025.pdf

5. Tennessee: the Tennessee Bureau of Investigation provides the actual number of background checks in their annual reports[62]

6. Utah: Utah's Bureau of Criminal Identification provides the actual number of background checks in their quarterly reports[63]

7. Virginia: Virginia State Police provides the actual number of background checks in their annual reports[64]

B.2 Details on market size and consumer demographics

From the 2019 5-year American Community Survey (ACS), we record the count of each zip code's population and its distribution by gender, age, education, race/ethnicity, citizenship, poverty, health insurance status, and marital status in the zip code.[65]

To measure zip code voting behavior, we download the precinct-level vote shares in the 2016 U.S. presidential election, from Voting and Election Science Team (2018).[66] We aggregate precinct-level vote shares for all conservative candidates in the presidential election in 2016 to construct a "Percent Conservative" measure for each precinct in the United States.[67] We then aggregate this to the zip-code level using the supplied Voting and Election Science Team (VEST) precinct shapefiles along with U.S. Census TIGER ZCTA shapefiles, weighting precincts by total number of votes.

We employ two measures of consumer engagement with the NRA. First, we use FEC data on large ($200+) individual contributions to the NRA's Political Action Committees from 2016-2022, tagged with the donor's zip code, which are mapped to the congressional district using the December 2018 crosswalk from the U.S. Department of Housing and Urban Development.[68] Second, as a complement, we use a dataset of grades issued to members of the U.S. House of Representatives in 2018 by the NRA, reflecting the political stances of locally elected politicians on gun rights legislation.[69]

[62] Source: https://www.tn.gov/tbi/divisions/cjis-division/firearm-background-checks.html

[63] No central repository exists but they can easily be searched for; See this for an example of 2020's https://bci.utah.gov/wp-content/uploads/sites/15/2021/01/2020Q4.pdf

[64] Source: https://vsp.virginia.gov/about-us/annual-report/

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

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

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

[68] Source: https://www.huduser.gov/portal/datasets/usps_crosswalk.html.

[69] Source: https://www.thetrace.org/2018/10/nra-grades-republican-candidates/. These are derived primarily based on voting history and a questionnaire the NRA issues to candidates; an example questionnaire can be found here: https://afj.org/wp-content/uploads/2020/01/Wilson-Attachments-p450-453.pdf

Our survey data on participation in the firearm market comes from Wave 26 of the American Trends Panel (ATP), conducted by Pew Research in April 2017. We define an individual's participation in the legal firearm market as a binary indicator equal to 1 if the respondent stated that they (i) "personally own any guns (NOT including air guns, such as paintball, BB or pellet guns)," or (ii) "could see myself owning a gun." We also used this survey to gather information on respondents' demographic and political characteristics, taking care to match these variables as closely as possible to the data definitions from the ACS and election returns.

We construct our summary of participation in the firearms market via LASSO-penalized linear regression on the survey data (Tibshirani 1996). Explicitly, and separately for each census region, we run a LASSO regression of our indicator for potential market participation on the full set of Pew-provided characteristics for gender, age category, education level, race/ethnicity, citizenship, marital status, income category, health insurance, and party affiliation. We use an adaptive LASSO procedure to then separately select the variables to include in each census region's linear probability model. Figure B.1 shows the coefficients from each census region's selected model.

We construct measures of adults in the firearm market by zip code-gender cells. To construct the share of adults in the market within each cell, we take the inner product between the vector of selected regression coefficients and each cell's demographic and political characteristics. We compute the market size of each cell as the product of our estimated share with the cell's observed adult population size, i.e.

$$M_{z,g} = P_z \cdot (\beta_r \cdot \bar{D}_{z,g}) \tag{B.1}$$

where z indexes zip-codes, g indexes gender, $P_{z,g}$ denotes the adult population according to the ACS, $\bar{D}_{z,g}$ denotes the average demographic in the zip-gender cell, and β_r denotes the (selected) LASSO regression coefficient for census region r.

B.3 Details on aggregate firearm purchase quantities

In this section, we describe adjustments made to publicly available NICS background check data to better reflect firearm purchases. For many firearm transactions since 1998, federal law requires the consumer first undergo a background check through the NICS system, called in by the firearm seller. To implement these background checks, the FBI administers a database that the party responsible for implementing the background check may query. The FBI makes counts of NICS queries to this database publicly available by state-month. These NICS counts have become a popular input for constructing measures of firearm purchases

(Lang 2016). The NICS dataset records the number of background checks, broken down by the weapon class purchased (handguns, long guns, or both types) and permit checks.

States differ in their NICS reporting on two important margins (Smucker et al. 2022). First, some states issue firearms permits that are exempt from reporting (known as "Brady exempt"), meaning that holders of these permits may purchase firearms without a NICS check, for up to five years since the permit was issued. This could lead to an underestimate of firearms via NICS checks. For some permits, these exemptions only apply to handguns. Second, some state agencies act as a Point of Contact ("POC") between the firearm seller and the FBI. The POC is a state agency (e.g., the New Jersey State Police) that takes responsibility for conducting background checks meeting both state and federal requirements. In implementing a background check for a single transaction, a state POC may query NICS multiple times. When this occurs, a measure of firearm purchases constructed from publicly available data on NICS queries may overstate the true number of firearm purchases that actually occur. State agencies differ in whether all background checks are conducted through a POC, or whether these POC checks are solely conducted on handgun purchases.

Figure B.5 displays the status of all US states in their reporting standards to NICS in 2016.[70] States in the South are more likely to have Brady exempt permits, suggesting that the NICS reporting may be biased downward in states with higher underlying demand for firearms. At the same time, states in the Pacific West are more likely to have POC checks on firearms, suggesting over-counting in this region.

Throughout, we rely on the standard adjustment to NICS checks from Brauer (2013), which adjusts background checks to account for multiple firearms purchased per background check. This formula is

$$NICS_{s,t} = 1.1 \cdot \text{HandgunChecks}_{s,t} + 1.1 \cdot \text{LongGunChecks}_{s,t} + 2 \cdot \text{MultipleFirearmChecks}_{s,t} \tag{B.2}$$

We follow a three-step approach to adjust NICS background checks to better reflect firearm purchases. First, in states where we observe physical transfers, we use these in lieu of NICS checks. Second, in states with POC checks, we estimate the relationship between NICS checks and POC checks to correct for multiple queries by state agencies. We convert these to purchases using Equation B.2. Third, in states with Brady exempt permits, we estimate a synthetic difference-in-difference model (Arkhangelsky et al. 2021) to adjust for under-reporting from Brady exemptions.

[70] Figure is copied from the 2016 NICS Operations report, downloaded from https://www.fbi.gov/file-repository/cjis/2016-nics-operations-report-final-5-3-2017.pdf/view

B.3.1 Over reporting in POC states

To correct for over-reporting from POC states, we manually collect POC information from state agencies for all states with POCs during our sample, where data was publicly available. Table B.1 displays the states with POCs and the data we were able to collect. For 11 of 17 states with a state POC, we are able to observe the state-level background checks to correct for multiple queries per transaction, and for 4 of these states, state agencies also report the number of physical transfers.

Figure B.6 displays the relationship between State POC background checks and physical transfers, for state-years in which we were able to construct both variables. The relationship, whether we allow for state-specific slopes or not, is very close to 1. We take this as evidence that POC background checks are a suitable proxy for firearm purchases, when permit/Brady exemptions do not apply.

We adjust observed NICS purchases in POC states to account for the fact that the state POC may query NICS more than once per background check. Specifically, we assume a constant adjustment factor

$$\text{StatePocChecks}_{s,t} = \alpha_s \cdot NICS_{s,t} \tag{B.3}$$

to convert NICS data from POC states to physical firearm transfers.[71]

Each state has a different constant of proportionality α_s, which we estimate as state fixed effects via OLS regressions between logged NICS and logged state POC background checks.

Table B.2 presents estimates of state-specific constants of proportionality α_s, along with an average estimator assuming no state heterogeneity. These state-specific constants have a low of 0.79 in Colorado, which implies each background check by the Colorado POC requires $1/0.79 \approx 1.27$ NICS checks. The high is 1.04 in Utah, which would be the case if some background checks from the Utah POC did not require a call to the NICS system. The average effect of $\alpha_{\text{average}} = 0.94$ implies that, for the average state-year, 100 POC checks require $100/0.94 \approx 106$ NICS checks. We apply state-specific adjustment factors to POC states where data on POC checks are available, and the average α_{average} to the remaining POC states. We then use the time series of adjusted NICS checks by state for the remaining correction for Brady exemptions.

The output of our process is an estimate of the number of NICS checks that would have occurred in each POC state in Table B.1, had they counterfactually not had their POC

[71] In states with partial permit POC, POC checks occur solely on handguns; as such, we convert only handgun counts via a state-specific constant, then apply Equation B.2 with the adjusted handgun count to produce our estimate of physical firearm transfers.

in place. These predictions are comparable to the number of NICS checks in state-years without a POC, and we refer to these comparable predictions as "NICS checks" in the following section.

B.3.2 Under reporting in Brady exempt states

To correct for under-reporting from Brady exemptions, we use historical changes in Brady exempt status by state, and construct a synthetic difference-in-difference estimator (Arkhangelsky et al. 2021) to estimate the effect of Brady exemptions on NICS background checks for each switching state. This approach uses 21 states (donors) with no Brady-exempt permits from 2004-2023 to construct a counterfactual time series of NICS checks in the absence of the introduction of a Brady-exempt permit. The synthetic difference-in-difference estimator differs from standard synthetic control in that it also allows for weights on time period to match pre-treatment trends. Explicitly, we assume the following data generating process for NICS checks per capita.

$$\log(\frac{NICS_{s,t}}{Pop_{s,t}}) = \vec{\theta_s} \cdot \upsilon_t + \gamma_s + \delta_t + \beta_s \cdot BradyExempt_{s,t} + \epsilon_{s,t} \tag{B.4}$$

This specification accommodates unobservable latent factors at both the time and state level. Table B.3 displays each treated state that moves to obtaining a Brady-exempt permit, along with the estimated effects of Brady-exempt status on total NICS checks (Equation B.2), handgun checks, and long gun checks. Brady exemptions lead to anywhere from a .6% decrease (Oklahoma) to a 49% decrease (Alabama) in NICS checks. The average is a 25% decrease. We apply the synthetic DiD state-specific estimates to the 10 states for which we observe a change in Brady-exempt status. For the other 13 states with Brady-exempt permits throughout our sample (i.e. no variation in treatment status), we apply the average effect of a 25% decrease in NICS checks.

We perform a similar exercise for the five states with Handgun permit exemptions (known as "Partial Permit POC" states). These states only exempt handgun purchases from NICS checks for those with valid permits and can demonstrate patterns of NICS records quite different from states with a simple Brady exemption. For these partial permit states, we use the transition from partial permit status to full reporting. Table B.4 displays the estimated effects, estimated via a synthetic difference in difference. Consistent with the exemption only applying to handguns, we estimate no effect on long gun checks, and a 180% *increase* in handgun checks from removing partial permit status. Iowa is the only Partial Permit POC state with changing status during our sample,[72] so we apply its state-specific adjustment to

[72]Michigan and New York ended their partial permit status prior to our sample, but we still use their

its handgun checks. In North Carolina and Nebraska, where we observe no change in Partial Permit POC status, we apply the 180% increase in handgun checks, averaged across the three switching states.

Figure B.7 displays the final ratio of the adjusted NICS checks to unadjusted NICS checks from 2016-2022, by state. On average, our adjustments increase purchases by 30% nationally, primarily due to the concentration of Brady exemptions in high demand states, such as the southern United States.

Our adjusted measure of NICS background checks simultaneously accounts for POC status, Brady exemptions, and Partial Permit POC rules. We use our adjusted measure of NICS background checks to construct moments in Section 5.3 as a more faithful estimate of firearm purchases by state-year across the United States.

B.4 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 which 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 compares these model identifiers in the same location of the strings. 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.[73] 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 that we do not ignore some transactions in our demand estimation simply due to a low quality string match. While this introduces estimation error to our demand system, the price-quantity results of Section 4 suggest that our matching approach captures the relevant variation in the data (e.g. demand slopes

estimates to impute the changes in Nebraska and North Carolina.

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

downward). 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 value in the FRB dataset, 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 across models.

State	Covered Firearms	Data	Background Checks	Approved Transactions	Physical Transfers	Gun Type
CA	All	Y	Y	Y	Y	Y
CO	All	Y	Y	Y	—	Y
D.C.	All	—				
FL	All	Y	Y	—	—	—
HI[1]	All	Y	—	Y	Y	Y
IL	All	Y	Y	—	—	—
MD[2]	Handgun	Y	Y	—	—	—
NH	Handgun	—				
NJ	All	—				
NV	All	—				
OR	All	Y	Y	—	—	—
PA	All	Y	Y	—	Y	Y
TN	All	Y	Y	—	Y	Y
UT	All	Y	Y	Y	—	—
VA	All	Y	Y	—	Y	Y
WA	Handgun	—				
WI	Handgun	—				

Table B.1: States with Point of Contact and Available Data

Notes: Left panel is a state's POC program. Right panel is the information recorded and stored by the program. A single approved transaction may involve multiple physical transfers of firearms.

[1]HI POC reports that background check practices confound the mapping between background checks, transactions, and physical firearm transfers.

[2]MD POC operates for handguns and "assault weapons;" however, most transactions of "assault weapons" have been banned in MD since October 1, 2013.

state	Coefficient	SE
california	0.8907	0.0294
colorado	0.7916	0.0232
florida	1.0083	0.0724
illinois	0.9366	0.0426
maryland	0.976	0.0444
oregon	0.9052	0.0411
pennsylvania	0.9973	0.0453
tennessee	0.9791	0.0287
utah	1.0397	0.0305
virginia	1.0276	0.041
Average	0.9416	0.0129

Table B.2: State-level Relationship between NICS Background Checks and POC Background Checks

Table displays estimates of the state-specific constants of proportionality α_s from Equation (B.3), used to correct for state POC agencies querying NICS multiple times per background check. For Maryland, whose POC covers only handguns, the converter is constructed from handgun checks only. Each coefficient is the exponentiated state fixed effect from an unweighted OLS regression of $\log(\text{POC checks}_{s,t}) - \log(\text{NICS checks}_{s,t})$ on state indicators (i.e., the within-state geometric mean ratio of POC checks to NICS checks), using a dataset of state-years with data on POC and NICS checks. Standard errors are computed via the delta method. The Average row reports the analogous estimate from a pooled regression on a constant only, which weights each state-year equally.

State	Brady Exempt Date	NICS Checks	Handgun Checks	Long Gun Checks
Michigan	2006-04	-0.391	-0.861	-0.178
		(0.122)	(0.548)	(0.104)
Kentucky	2006-07	-0.408	-0.563	-0.336
		(0.183)	(0.540)	(0.165)
Kansas	2011-05	-0.200	-0.320	-0.151
		(0.151)	(0.176)	(0.171)
Nevada	2011-09	-0.347	-0.379	-0.418
		(0.121)	(0.118)	(0.168)
West Virginia	2014-06	-0.180	-0.216	-0.154
		(0.111)	(0.369)	(0.271)
Louisiana	2015-04	-0.136	-0.138	-0.151
		(0.128)	(0.318)	(0.240)
Alabama	2016-03	-0.683	-0.760	-0.593
		(0.128)	(0.256)	(0.151)
South Dakota	2017-03	-0.223	-0.181	-0.215
		(0.110)	(0.150)	(0.124)
Minnesota	2017-05	-0.260	-0.397	-0.153
		(0.064)	(0.084)	(0.083)
Oklahoma	2023-06	-0.008	-0.002	-0.063
		(0.107)	(0.165)	(0.108)
Average		-0.296	-0.400	-0.258
		(0.062)	(0.085)	(0.055)

Table B.3: Estimated Effects of Introducing Brady Exempt Permits on NICS Background Checks

Table displays the estimated effects (in log points) of introducing a Brady exempt permit on background checks, estimated separately for each of the 10 states we observe switching to Brady exempt status from 2004-2023, via the synthetic difference-in-differences estimator of Arkhangelsky et al. (2021). Outcomes are, respectively, the log of per-capita NICS checks (Equation B.2), per-capita handgun checks, and per-capita long gun checks. All check counts are first adjusted for POC over-reporting using the state-specific factors of Table B.2 and Equation (B.3). Brady Exempt Date is the first month in which the state's exemption is in effect. Standard errors in parentheses are computed via the placebo variance estimator of Arkhangelsky et al. (2021). The Average row weights each state-specific estimate by its number of post-treatment months, and computes a standard error via a nonparametric bootstrap with 100 replicates, using stratified resampling of states from the pools of treated and donor states.

State	Partial Permit POC End	Handgun Checks	Long Gun Checks
New York	2004-09	0.194	0.080
		(0.069)	(0.060)
Michigan	2012-12	2.009	-0.164
		(0.360)	(0.156)
Iowa	2021-07	2.686	0.104
		(0.103)	(0.079)
Average		1.836	-0.082
		(0.087)	(0.056)

Table B.4: Estimated Effects of Ending Partial Permit POC Status on NICS Background Checks

Table displays estimated effects (in log points) of ending Partial Permit POC status (i.e., moving handgun purchases by permit holders back to NICS reporting) on background checks, estimated separately for each of the three states we observe ending partial permit status (New York in 2004, Michigan in 2012, and Iowa in 2021), via the synthetic difference-in-differences estimator of Arkhangelsky et al. (2021). Outcomes are the log of per-capita handgun checks and per-capita long gun checks. All check counts are first adjusted for POC over-reporting using the state-specific factors of Table B.2 and Equation (B.3). Partial Permit POC End is the first month of full NICS reporting. Standard errors in parentheses are computed via the placebo variance estimator of Arkhangelsky et al. (2021). The Average row weights each state-specific estimate by its number of post-treatment months, and computes a standard error via a nonparametric bootstrap with 100 replicates, using stratified resampling of states from the pools of treated and donor states.

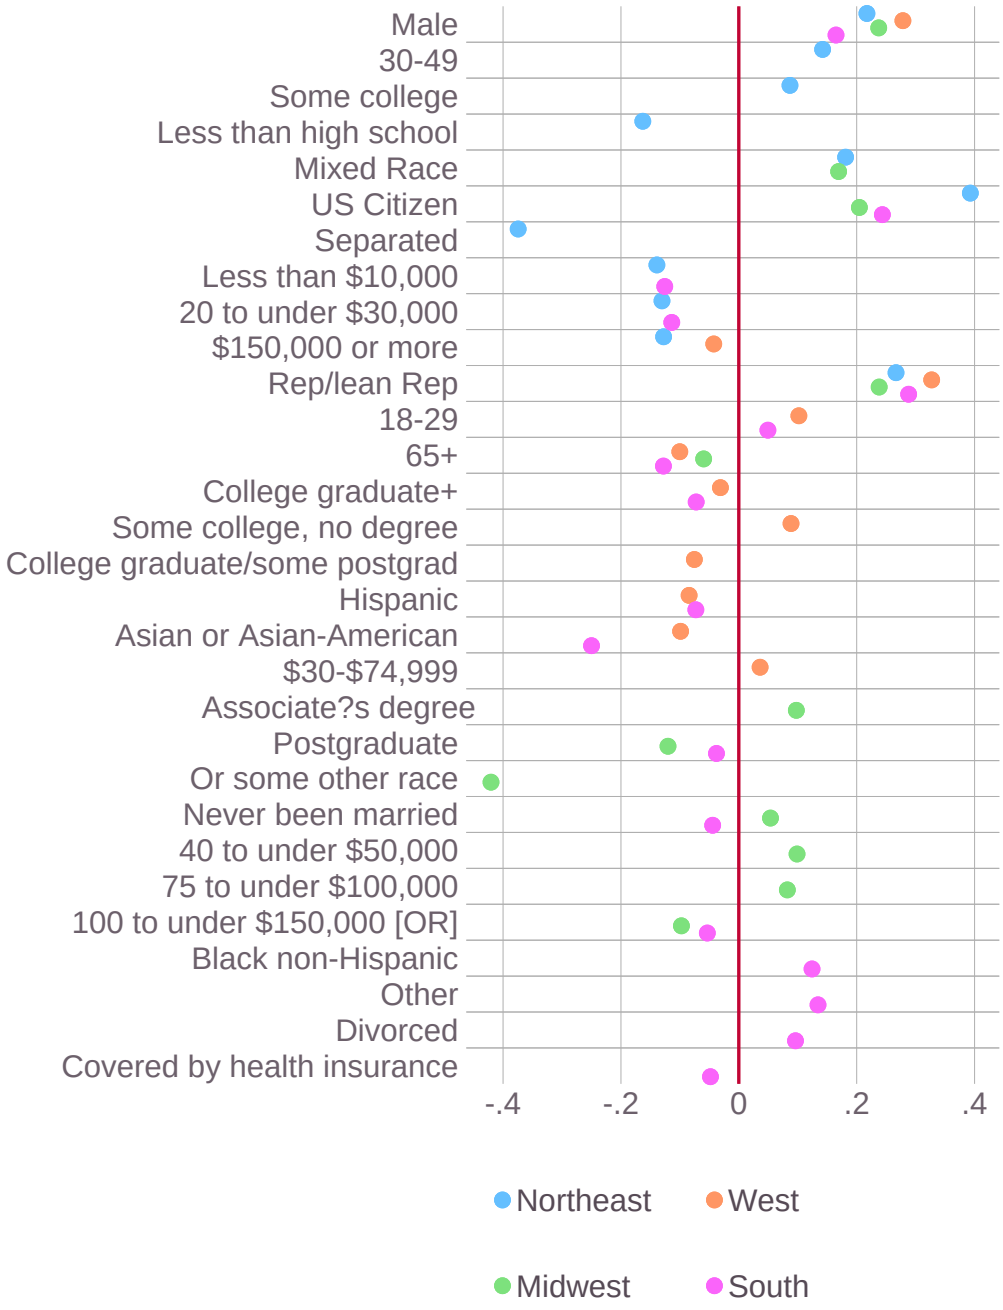
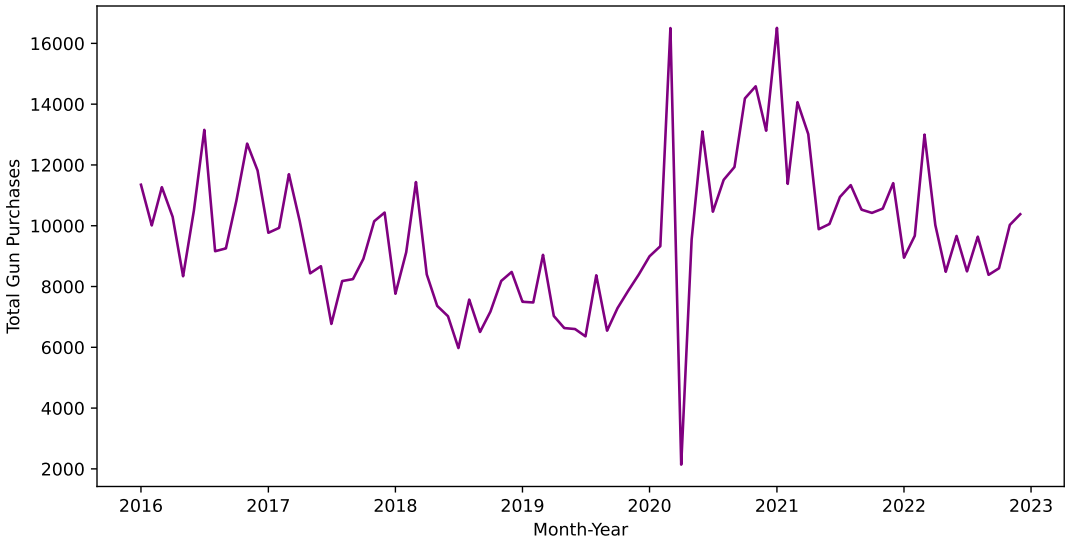
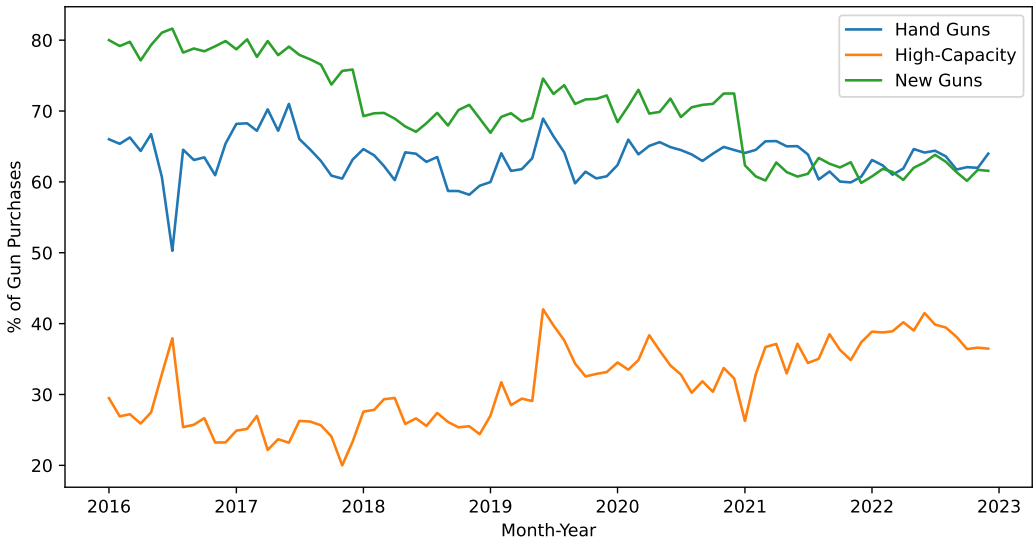


Figure B.1: Estimated Lasso Coefficients for Predicting Market Size, From Pew's American Trends Panel

Figure displays the estimated coefficients from LASSO-penalized regressions predicting participation in the firearm market, estimated separately for each census region on respondents to Wave 26 of Pew's American Trends Panel. The outcome is an indicator equal to one if the respondent personally owns a gun or reports they could see themselves owning one. Candidate predictors are the full set of survey indicators for gender, age category, education level, race/ethnicity, citizenship, marital status, household income category, health insurance coverage, and political party affiliation. Each region's linear probability model is selected via an adaptive LASSO with penalties chosen by 10-fold cross-validation. Only predictors selected (non-zero) in at least one region are displayed, and only selected coefficients are plotted.



(a) Aggregate Gun Purchases Over Time



(b) Composition of Purchases Over Time

Figure B.2: Firearm Purchase Trends in Massachusetts

Figure displays the time series of gun purchases in Massachusetts from 2016-2022. Panel (a) displays the total number of gun purchases each month, while Panel (b) displays the composition of these purchases across different types of firearms. Handguns refers to those transactions classified as a handgun in the FRB data, and the percentage on the vertical axis uses total firearm transactions each month as a denominator. For new gun and high-capacity firearms—these characteristics are only defined for purchases we match to the blue book data—and the percentage on the vertical axis uses total firearm transactions linked to an entry in the BBGV each month as a denominator.

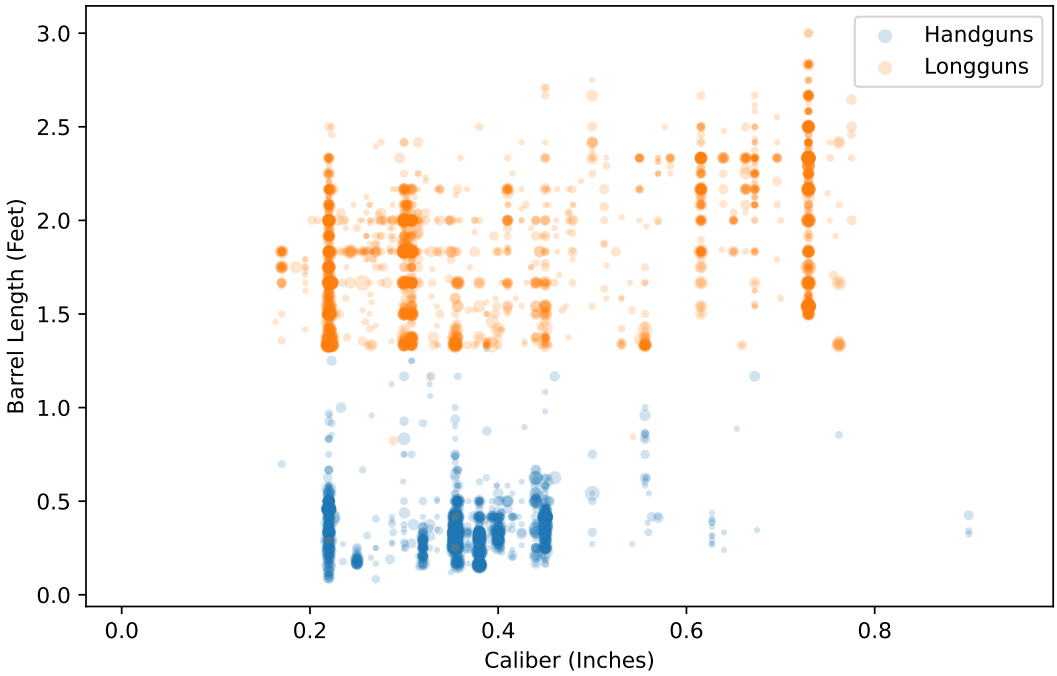


Figure B.3: Caliber and Barrel Length Distribution Across Gun models

Figure displays the distribution of gun models by caliber (in inches) and barrel length (in inches), split by weapon class. Each dot is a firearm model, with size proportional to the number of purchases in Massachusetts during our sample.

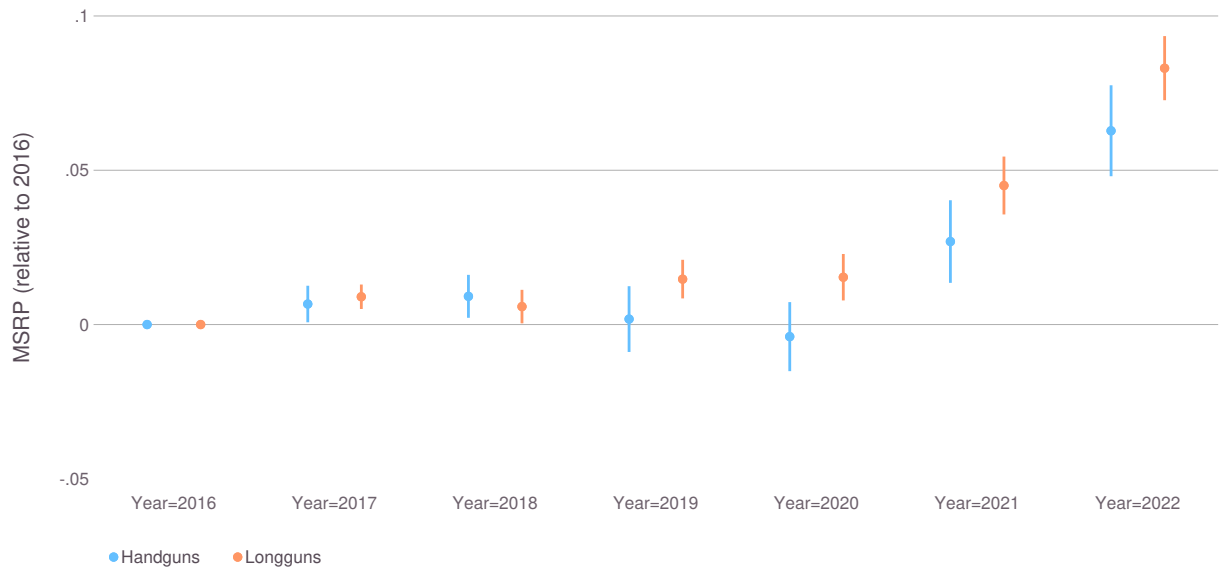


Figure B.4: New Gun Prices over Time (Controlling for Variety)

Figure displays estimates from an unweighted regression of log(MSRP) onto fixed effects for firearm model j year t, separately by weapon class. The sample of firearm models used are those (i) under active production and (ii) matched to at least one firearm purchase in Massachusetts. Points and 95% confidence intervals are for the year fixed effects, with 2016 the excluded category for each class. Standard errors are clustered by firearm model.

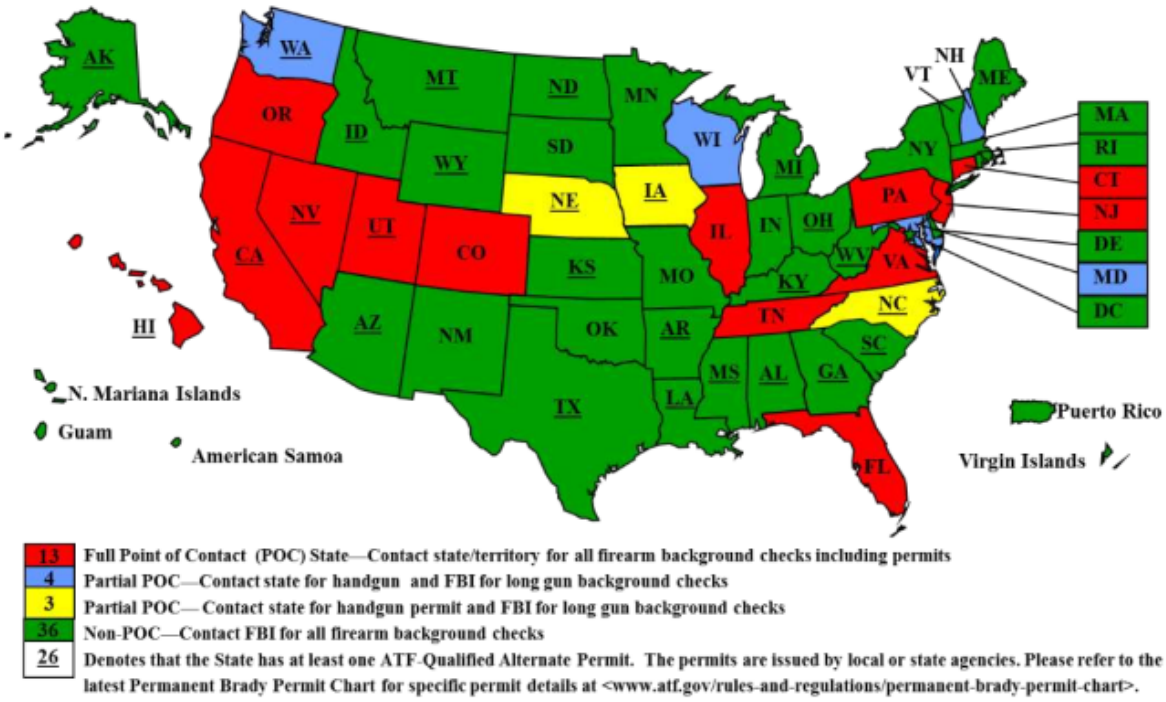


Figure B.5: POC and Brady Exemptions by State, 2016

Figure shows an example POC Status report from the FBI's 2016 NICS Operations Report. Colors denote various POC designations, while underlined state names denote Brady Exempt Status. Figure is sourced from https://www.fbi.gov/file-repository/cjis/2016-nics-operations-report-final-5-3-2017.pdf/view.

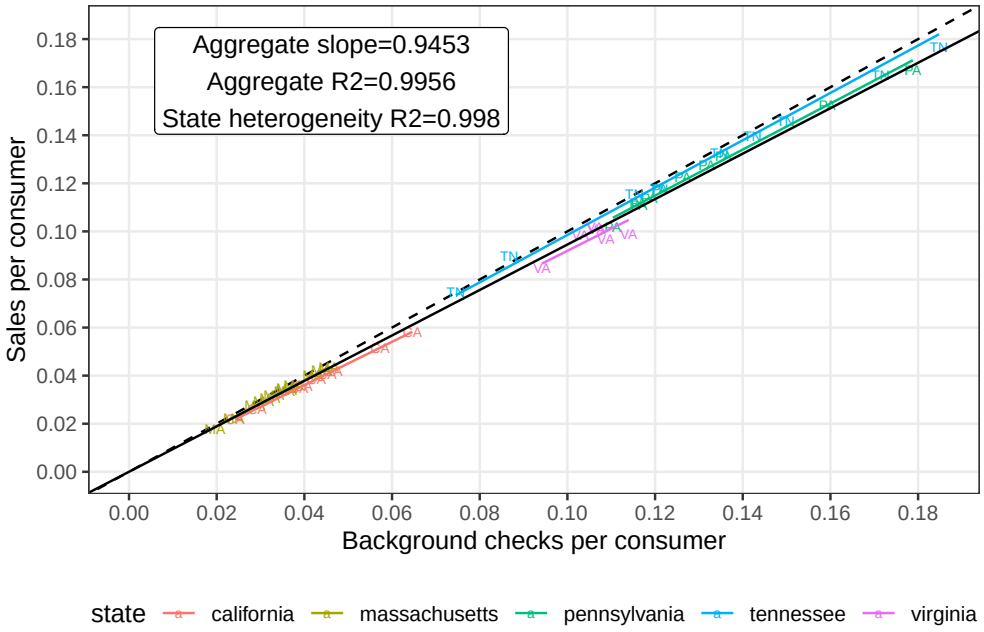


Figure B.6: Relationship between POC Background Checks and Physical Firearm Transfers

Notes: Figure shows the relationship between firearm sales per consumer against Background checks per consumer, a sample of state-year data from 2010 onward. We observe physical firearm transfers in CA, MA, PA, TN, and VA. For all states except MA, we report state POC background checks. In MA, we report NICS background checks since MA does not have a state POC. Although we observe physical transfers in HI, we do not reliably observe POC checks in that state. All regression statistics are computed in logs. Aggregate slope denotes the OLS fit from a pooled regression across states. Aggregate R2 denotes the R^2 from this regression. State Heterogeneity R2 denotes the R^2 from a regression allowing state specific slopes.

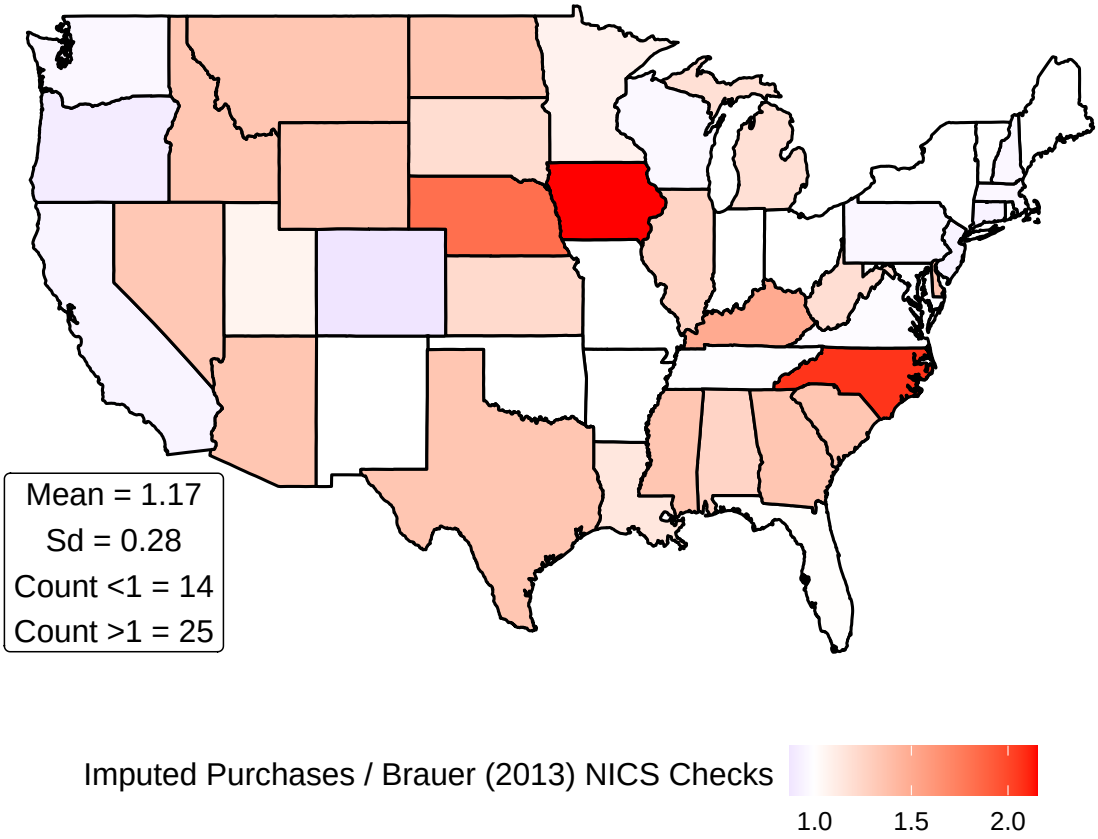


Figure B.7: Ratio of Adjusted NICS Background Checks to Unadjusted NICS Background Checks, 2016-2022, by State

Figure displays, for each state, the ratio of our final imputed firearm purchases to the Brauer (2013) aggregation of NICS background checks, averaged over 2016-2022. The denominator is the unadjusted NICS check count of Equation B.2, developed by Brauer (2013). The numerator is our imputed purchase series: observed physical firearm transfers in state-years where available, and otherwise NICS checks adjusted for POC over-reporting (Table B.2), partial permit exemptions (Table B.4), and Brady exemptions (Table B.3). White denotes a ratio of one: no net adjustment across years, relative to the measure of Brauer (2013). Reported summary statistics are unweighted across the mapped states. States with a ratio of exactly one are included in neither count.

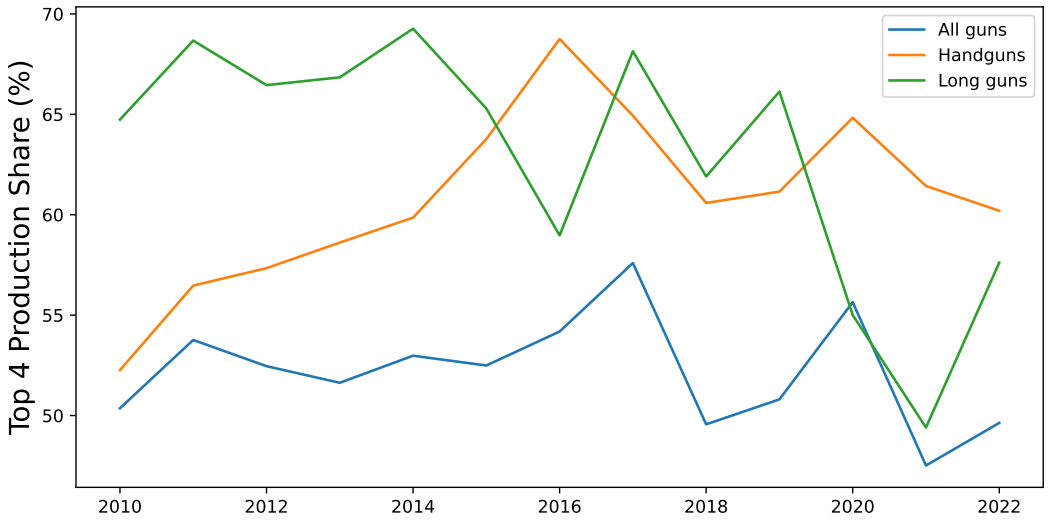


Figure B.8: Concentration in the U.S. firearms industry

Figure displays the yearly share of the four largest firms by their reported firearm production-and-sale into the domestic consumer firearm market. Series are the yearly top-four share across all firearms, and separately by handguns and long guns. Data are derived from the ATF Annual Production data (by plant), and merged across plants owned by the same firm using our hand-collected data on the ownership structure of firearm manufacturers.

C Diagnostics on Cost Shock Instrument

This section describes robustness checks and supporting details for the cost shock instruments, constructed in 4.1, and used to estimate both the price elasticity regressions in Section 4.2 and the equilibrium model in Section 5.

C.1 Robust to the granularity of PPI commodities in $\mathcal{K}$

While our preferred specification uses the most granular PPI commodities possible (6-digits), Table C.1 applies the same IV-LASSO estimator of Equation (4), using higher-level aggregations of the PPI commodities. Specifically, we estimate the equation using 2-digit, 4-digit, and (our preferred) 6-digit aggregations, which yield 11, 33, and 42 unique commodity time series, respectively. While the potential instrument set and the number of instruments chosen changes slightly (5 at lower digits), the estimated price elasticity hardly changes, and all specifications have a highly predictive first stage.

C.2 Robust to alternate penalty schemes for model selection

The 7 selected instruments in our preferred specification are chosen in a data-driven manner, specifically developed by Belloni et al. (2012) to balance a bias-variance tradeoff and deliver asymptotically valid standard errors when plugged into 2SLS. Nonetheless, it is reasonable to contemplate other model selection criteria that could imply a different LASSO penalty parameter and select a different set of instruments. We assess robustness to the model-selection criteria, by re-solving the IV-LASSO estimator along a grid of penalty parameters, each selecting a potentially different quantity and composition of instruments.

Figure C.3 plots the estimated price elasticities, as a function of the number of selected instruments. The Figure documents that among criteria selecting 1–14 instruments, the estimated price elasticity is stable around -2.1. After this point, the introduction of additional instruments attenuates the estimated price elasticity coefficient towards zero, exactly the weak instruments problem the IV-LASSO routine is designed to address. Still, in the neighborhood of our preferred specification with 7 selected instruments, selecting slightly more or slightly fewer commodity-manufacturer interactions does not substantially alter our estimated price elasticity of firearm demand.

C.3 No anticipatory effects, rejecting unobserved pre-trends

Another natural concern with our instrument is that the manufacturer-specific cost index $\sum_k \hat{\beta}_{m,k} \log(\tilde{p}_{k,t})$ does not capture the genuine passthrough of input prices, but rather proxies

for a confounding, manufacturer-specific pre-trend of price changes. To formally test this possibility, we run the following regression:

$$\log(p_{j,t}) = \tilde{\gamma}_j + \tilde{\eta}_{c(j),t} + \beta_0(\sum_{k\in\mathcal{K}} \hat{\beta}_{m(j),k} \log(\tilde{p}_{k,t})) + \beta_1(\sum_{k\in\mathcal{K}} \hat{\beta}_{m(j),k} \log(\tilde{p}_{k,t+1})) + \epsilon_{i,j,t}, \quad \text{(C.1)}$$

where imposing $\beta_0 = 1, \beta_1 = 0$ reduces the above equation to our preferred first-stage specification. The regressor associated with β_0 is equivalent to the "prediction" of log(MSRP) from only the cost shocks. In this exercise, we hold the terms in these cost shocks fixed at the values from our preferred specification in Section 4.1.

The new term in Equation (C.1) is the regressor associated with β_1: this is the prediction of log MSRP if gun j was exposed to cost shocks that occurred *next* year. If our estimates rely on variation from genuine shocks to input prices, there should be no residual relationship between MSRP today and cost shocks that realize next year. However, if our selected instruments simply proxy for pre-trends in pricing behavior at the selected manufacturers, we would expect a positive relationship between cost shocks both today and in the future.

Table C.2 displays estimates of Equation (C.1) on the subsample of manufacturers with non-zero coefficients (i.e. the 4 manufacturers for which the IV-LASSO procedure recovers a coefficient $\hat{\beta}_{m,k} \neq 0$). By construction, when we include only the MSRP predicted from the input prices in period t, the coefficient in Column (1) is $\beta_0 \approx 1$. While, Column (2) tests whether the input prices in period $t+1$ have any predictive power for the MSRP in period t (conditional on input prices in t).[74] We estimate $\beta_1 \approx 0$ with good precision, suggesting that there is no anticipatory effect of future cost shocks on current MSRP, and that our estimates are not based on manufacturer-specific pre-trends in MSRP that happen to be proxied by our selected cost indices.

C.4 Similar predictive performance for out-of-sample firearms

We evaluate the degree to which our estimated relationship between the prices of inputs (commodity metals) and outputs (firearms) is a structural feature of each manufacturer. To do this, we evaluate the predictive power of our cost shocks, $\widehat{\log(p_{j,t})} = \sum_{k\in\mathcal{K}} \hat{\beta}_{m(j),k} \log(\tilde{p}_{k,t})$, on prices using data on firearm MSRP from two separate sub-samples not seen by our LASSO estimator.

The first sub-sample comprises firearms models outside our estimation sample in Section 4.1, that are made by the same manufacturers in our sample during the same period (2016–

[74] Because input prices are serially correlated, we continue to control for the MSRP predicted by contemporaneous cost shocks when testing the relationship between observed MSRP in period t and the MSRP predicted from cost shocks in period $t+1$.

2022). As discussed in Section 4.1, we restrict our baseline estimates for the price elasticity of demand to those models for which at least 100 purchases are made in Massachusetts from 2016-2022. This filtering is done primarily to avoid selection bias in the log transformation of purchases, as firearms with zero purchases in Massachusetts drop out of Equation (5), which excludes "niche" models that are less popular. This bias only affects the second stage of estimation. If the manufacturer-specific coefficients $\hat{\beta}_{m,k}$ we recover are features of manufacturers, they should be valid for predicting the prices of both the popular and "niche" models made by the same manufacturer.

The second sub-sample comprises the same firearm models from our estimation sample in Section 4.1 during the years *before* our estimation sample (2010-2015). While manufacturers may change their production practices and input mix over time, if the input dependencies we recover are structural features of each manufacturer, MSRP should be linked to prices of the same inputs in the six years prior to our estimation sample.

Figure C.4 plots the binscatter of the relationship between input prices and observed MSRP of each sub-sample, for manufacturers with at least one non-zero pass-through parameter $\hat{\beta}_{m,k}$. As in Table C.2, we recover a coefficient near 1 (0.993) on predicted MSRP for our main estimation sample. The alternative sub-samples—comprising "niche" firearms during our estimation sample, and popular firearms before our estimation sample—yield coefficients of 0.910 and 1.113, respectively. These coefficients indicate that the cost shocks we construct are similarly predictive of firearm prices, both in- and out-of-sample. Moreover, the binscatter shows that this relationship is approximately linear in all samples. These results suggest that the estimated relationships between the input prices of metal commodities and output prices of firearms are structural features of the selected manufacturer.

A notable difference across the sub-samples is that the pre-period coefficient has a large standard error (0.507). This is attributable to the fact that there is relatively little price variation from 2010-2015, since the supply chain shocks attributable to COVID-19 do not occur over that period, limiting variation in the independent variables.

C.5 Patent Validation

Our preferred approach to constructing cost shocks selects 7 instruments from a set of hundreds of manufacturer-metal combinations. Although the IV-LASSO procedure of Belloni et al. (2012) is designed specifically to address the challenge of selecting a useful subset among many weak instruments, one may still worry that our chosen instruments capture spurious correlations of price movements that have little to do with actual underlying cost structures at a manufacturer. Absent detailed production data on the metal composition

used to produce different firearms, there are few public resources available to validate our chosen instruments.

We rely instead on patents data, and compare the metals associated with each selection of a manufacturer-commodity pair, to the prevalence of those metals in that manufacturer's firearm patent filings. The idea is, if the selected instruments genuinely reflect the manufacturer's input mix, rather than a spurious correlate, that manufacturer should disproportionately mention the metal in its own engineering documents. Patents provide the only large-scale, manufacturer-attributable, machine-readable description of firearm construction that is publicly available, so we use the text of each manufacturer's patents as a proxy for its material inputs. Although imperfect as a measure of production practices, we are able to test whether the metal inputs implied by the selected instrument are over-represented in that manufacturer's patent filings, relative to its competitors'.

To implement this test, we begin by scraping all patents filed in 2000 or later for all 19 manufacturers with at least 15 model-year observations in the sample used to generate our IV-LASSO results (Table 3), from the Google Patents BigQuery public dataset.[75] These 19 manufacturers constitute 91% of all new firearms purchases in Massachusetts, and 75% of all model-year observations. We exclude manufacturers that are wholly-owned subsidiaries, who file their patents under their parent company name (e.g., Stoeger, which is owned by Beretta, has zero patents by construction). We restrict attention to patent families (the de-duplicated ID of a patent for multiple filings) with (i) at least one filing in the United States and (ii) a Cooperative Patent Classification (CPC) code beginning with `F41A` (functional features of firearms) or `F41C` (small arms). This ensures we do not double count patents that are re-issued/filed in multiple countries. When available, we analyze the text of the granted patent, but if unavailable, we use the patent application. The document used to represent each patent consists of the description, the abstract, and the claims. This yields a dataset with the description of 895 unique inventions attributable to 19 manufacturers.

Within our patent dataset, we then identify references to manufacturing practices that use various materials. We locate references in each patent text to 14 material groups: steel, aluminum, iron, copper, brass, nickel, titanium, zinc, tungsten, polymer (e.g. plastics), rubber, wood, carbon fiber, and precious metal (e.g., gold and silver). These match the material categories present in our PPI cost indices, along with other common inputs to firearms. We flag each sentence that mentions a material, and count it as indicative of a material being part of the production process if the sentence also contains the words "made", "fabricated", "formed", "constructed", or "manufactured", with these words followed by

[75] Initial queries of smaller manufacturers in our data yielded no results for patents, and none of our IV-LASSO specifications yielded non-zero coefficients for these smaller manufacturers.

"of", "from", or "with" (e.g., "made of", "constructed with"). We also flag sentences with the phrases "preferably", and "in a preferred embodiment". We discard all sentences with the phrases "in contrast to", "instead of", and "rather than", in order to count only sentences that actually suggest the material is used, as opposed to sentences that reject a particular material.[76]

We then aggregate across patents by each manufacturer to construct a manufacturer-by-material dataset, counting mentions of each material used in manufacturing. On average, each patent contains 2.46 mentions (sentences) that reference a material group explicitly used in production. Though, about half of the patents have no mention of any material group (likely referring to other parts of the production process, such as streamlining manufacturing or organization, or mechanisms that are material-invariant). On average, each manufacturer has 116 mentions of a material across their patent portfolios.

Our final step maps the materials mentioned in manufacturing patents to the pass-through parameters we estimate in Section 4.1. We first hand-code the 42 metal commodities measured in the PPI to material groups: iron, steel, aluminum, copper, titanium, tungsten, zinc, brass, and nickel.[77] We then count the mentions of each relevant material group in patents, among the manufacturers with non-zero pass-through parameters on the relevant PPI commodity. For interpretation, we divide this by the total number of material mentions across each manufacturer's patent portfolio to construct a "material intensity index": the share of each manufacturer's material mentions that reference the relevant metal.

To evaluate whether the prevalence of each material aligns with the chosen commodity coefficients for each manufacturer, we compare their material intensity index to the average index across "control" manufacturers: the 15 manufacturers in our primary specification that have 0 selected IV-LASSO coefficients.[78] The logic of this comparison is that, conditional on fixed effects for gun model and class-year, the coefficients selected by the IV-LASSO $\beta_{m,k}$ represent a manufacturer's exposure to a commodity *relative* to the average. For example, a negative coefficient on steel does not imply the manufacturer uses negative amounts of steel in firearm production, but rather that they are exposed to steel prices in a detectably lower way than a typical manufacturer (i.e., steel is less important for the selected manufacturer). We compute the difference in a metal's patent mentions between the "treated" manufacturer

[76] For example, a sentence with the phrase "made of an aluminum frame, instead of steel" would count for neither material group, since we cannot systematically distinguish whether aluminum or steel is the material actually used without hand inspection.

[77] In some aggregate categories, even at the 6-digit level, the PPI includes multiple metals. For example, PPI 10-220181, "Other non-ferrous metals", includes titanium, nickel, tungsten, and zinc (U.S. Bureau of Economic Analysis 2017).

[78] This statistic is an unweighted average across manufacturers. Weighting by the count of patents/mention hardly changes the control averages.

that has a non-zero pass-through parameter and the control mean, which we then compare to the recovered coefficient from our IV-LASSO first stage.

Table C.3 displays the results for the 7 selected coefficients in our preferred specification (6-digit, class-year fixed effects).[79] Across the board, relative material prevalence is consistent with the sign and magnitude of the selected PPI coefficients. Two of Sturm Ruger's coefficients correspond to steel materials, and correspondingly, they have nearly double the prevalence of steel in their patents (31.3% versus 16.7%), aligning with the large positive coefficients. The other coefficient for Sturm Ruger is also positive, and corresponds to unalloyed nonferrous metals (a catch all for rarer metals) which similarly is 4 times more prevalent in Sturm Ruger's patents. Both Smith & Wesson and Mossberg have positive coefficients on aluminum, and have excess prevalence of aluminum in their patents. Though, the difference is admittedly small for Mossberg, possibly attributable to an overall low number of patents/material mentions. Finally, Glock has *negative* loadings on both hot-rolled steel and aluminum, consistent with their dependence on polymer (plastics) in the manufacturing of their weapons (e.g. metals overall play a smaller role in the production of Glock weapons).

Our quantitative results are consistent with anecdotal evidence linking these firearm manufacturers to their associated metals. Sturm Ruger explicitly lists steel prices as a risk factor in their 2020 10-K filing with the SEC[80]. Smith & Wesson transitioned some of its most popular models to a proprietary scandium-aluminum alloy in 2001.[81] Glock's initial key differentiator in the 1980's was its reliance on polymer/plastics.[82] Our patent analysis reveals that the selected instruments are not simply spurious correlates: they reflect underlying cost structures and production practice of manufacturers, present in their patent filing and business strategy.

[79] Across specifications with different granularity (e.g. 2-, 4-, and 6-digit PPI commodities), only a single coefficient is sign-inconsistent with relative material prevalence in patents filings. That is, 16/17 coefficients align, and the sign reversal is numerically small in magnitude.

[80] Source: https://ruger.com/corporate/PDF/10K-2020.pdf

[81] Source: https://gunsmagazine.com/guns/the-j-frame-goes-scandium/

[82] Source: https://www.npr.org/2012/01/24/145640473/how-the-glock-became-americas-weapon-of-choice

	(1) 2-digit	(2) 4-digit	(3) 6-digit
Log(MSRP)	-2.263**	-2.113**	-2.136**
	(0.880)	(0.852)	(0.831)
Observations	2436	2436	2436
Gun Model FE	Yes	Yes	Yes
Year FE	Class x Year	Class x Year	Class x Year
# Potential Instruments	726	2178	2772
# Selected Instruments	5	5	7
Sup-Score Test Statistic	12.48	12.57	12.57
First Stage Sig. (p-value)	0.0000	0.0000	0.0000
Method	IVLASSO	IVLASSO	IVLASSO

Table C.1: Estimated Price Elasticity using Different Aggregations of commodities

Table displays the estimated own-price elasticity of demand from the IVLASSO routine described in Section 4.1 applied to a sample of 2,436 firearm model-years. The sample is restricted to firearm models with at least 100 purchases in Massachusetts between 2016–2022. Dependent variable is the log of a firearm model's purchase quantity in Massachusetts each year. Regressor is the log of a firearm's MSRP each year. Columns vary in the granularity of the price indices $\tilde{p}_{k,t}$ used: (1) 2-digit commodities, generating 11 indices; (2) 4-digit commodities, generating 33 indices; and (3) 6-digit commodities, generating 42 indices, our preferred specification from Table 3. All columns include fixed effects for gun model j and class-year $c(j), t$. #Potential Instruments is the number of manufacturer-commodity pairs assumed to be valid cost-shock instruments, and #Selected Instruments is the number of these manufacturer-commodity pairs chosen by the IVLASSO procedure to optimize estimation performance. Sup-Score Test Statistic is the statistic from a joint test against the irrelevance of all the instruments selected by the IVLASSO procedure, with the associated p-value below (i.e., low p rejects the null of irrelevant instruments, adjusting for the LASSO selection procedure). Standard errors robust to heteroskedasticity in parentheses. * denotes $p < .1$, ** denotes $p < .05$, *** denotes $p < .01$.

	(1) Log(MSRP)	(2) Log(MSRP)
LASSO Cost Shock	0.994***	0.973***
	(0.032)	(0.044)
Cost Shock Next Year		0.028
		(0.042)
Observations	929	929
Within R-Squared	0.46	0.46
Gun Model FE	Yes	Yes
Class-Year FE	Yes	Yes

Table C.2: Anticipatory Effects of Cost Shock

Table displays regression estimates of Equation (C.1), using the contemporaneous realization and lead of the LASSO cost shock instrument based on variation in PPI metal commodities. The sample is restricted to firearm models that (i) have at least 100 purchases in Massachusetts between 2016–2022 and (ii) are produced by Sturm Ruger, Smith & Wesson, Glock, or Mossberg: the four manufacturers in Figure C.2 for which our procedure selects non-zero manufacturer-commodity coefficients $\beta_{m,k}$. LASSO Cost Shock denotes the predicted MSRP for period t, based only on the contemporaneous (period t) input prices $\sum_{k\in\mathcal{K}} \hat{\beta}_{m(j),k} \log(\tilde{p}_{k,t})$. Cost Shock Next Year denotes the predicted MSRP for period t, based only on the leading (period $t+1$) input prices $\sum_{k\in\mathcal{K}} \hat{\beta}_{m(j),k} \log(\tilde{p}_{k,t+1})$. Within R^2 denotes the percent of variation explained by the cost shocks after residualizing the class-year and model fixed effects. All columns include fixed effects for gun model j and class-year $c(j),t$. Standard errors robust to heteroskedasticity in parentheses. * denotes $p < .1$, ** denotes $p < .05$, *** denotes $p < .01$.

Manufacturer	# Patents	# Mat Sents.	PPI	Rel. Metal	Mfr. %	Ctrl. %	Δ%	$\hat{\beta}$
Sturm, Ruger & Co.	86	638	Alloy steel castings	Steel	31.3	16.7	14.6	+0.143
Sturm, Ruger & Co.	86	638	Cold-rolled steel	Steel	31.3	16.7	14.6	+0.145
Sturm, Ruger & Co.	86	638	Unalloyed nonferrous	Ti/Ni/W/Zn	17.4	4.4	13.0	+0.049
Smith & Wesson	152	231	Aluminum scrap	Aluminum	21.2	11.4	9.8	+0.073
Mossberg	19	29	Aluminum scrap	Aluminum	13.8	11.4	2.4	+0.100
Glock	36	14	Hot-rolled steel	Steel	7.1	16.7	-9.6	-0.056
Glock	36	14	Aluminum die-castings	Aluminum	7.1	11.4	-4.3	-0.056

Table C.3: Patent-text validation of selected cost-shock instruments

Table uses the procedure of Section C.5 to show each manufacturer-commodity pass-through coefficient selected by the LASSO first stage (Equation (4), at the 6-digit commodity), in comparison to the prevalence of the corresponding metal in that manufacturer's patent filings. Rows are the selected manufacturer-commodity pairs. Columns are different features of each pair. *# Patents* is the number of unique post-2000 firearm patents (after family de-duplication) attributed to the manufacturer. *# Mat Sents.* is the manufacturer's total count of sentences in patent text for which we identify a manufacturing practice based on a material. *PPI* is the selected commodity price series. *Rel. Metal* is the raw metal(s) priced by the selected series. *Mfr. %* is the share of the manufacturer's *# Mat Sents.* that reference the *Rel. Metal* in the selected *PPI*. *Ctrl. %* is the average value of this statistic, equally weighted, across all patent-holding manufacturers that have no pass-through coefficient selected by the LASSO first stage. Δ% is the difference *Mfr. % − Ctrl. %*. $\hat{\beta}$ is the value of the pass-through parameter $\beta_{m,k}$ estimated from the LASSO first stage. The instrument is sign-consistent for a row when $\text{sign}(\Delta\%) = \text{sign}(\hat{\beta})$, which holds for every selected manufacturer-commodity pair in our preferred specification.

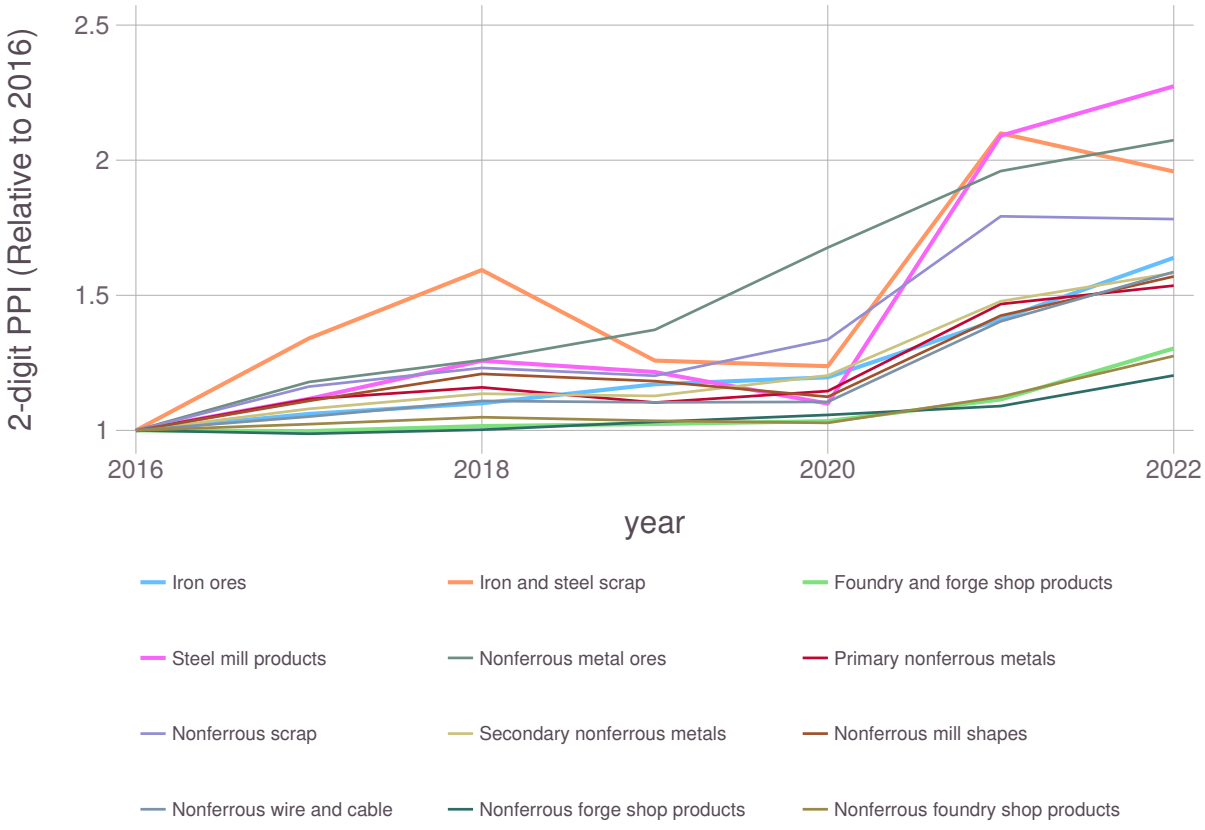


Figure C.1: Time Series of Annual Price Indices for Firearms Production Inputs

Figure shows the annual time series of all 2-digit primary metal commodity PPIs (those beginning with 10), from the Bureau of Labor Statistics. Each PPI is normalized to be 1 in 2016.

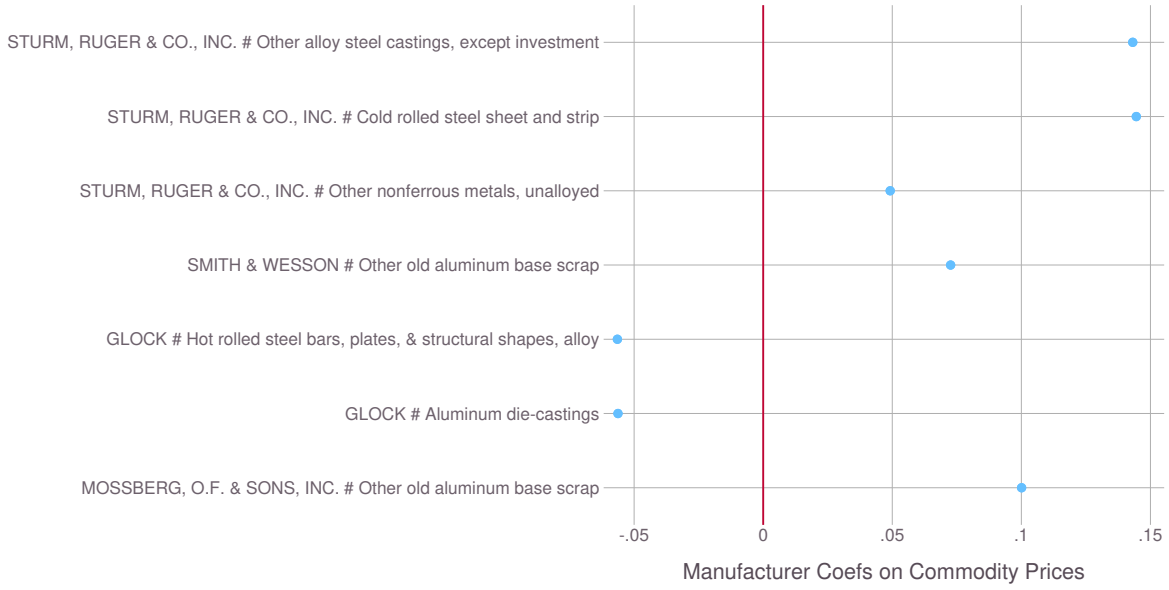


Figure C.2: Estimated Non-Zero coefficients for IVLASSO of Prices

Note: Figure plots the values of the selected coefficients from a LASSO regression of log(MSRP) on the interaction between log commodity price indices and manufacturer fixed effects, further conditioning on fixed effects for gun model j and class-year c, t, as in Equation (4). Estimation is on a sample of firearm model-years, restricted to firearm models with at least 100 purchases in Massachusetts between 2016–2022.

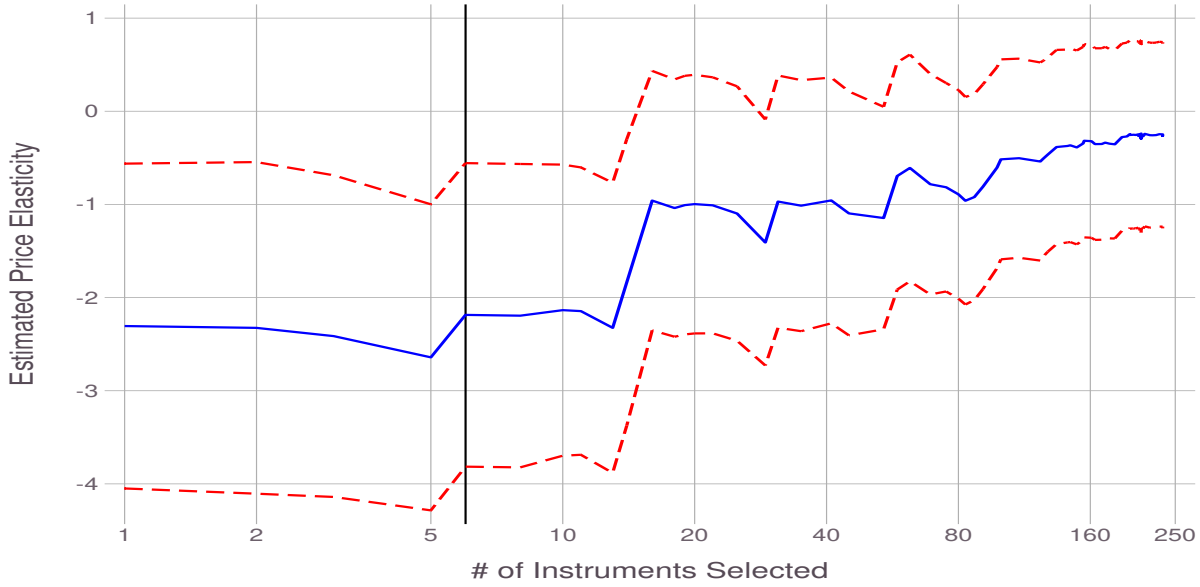


Figure C.3: IVLASSO Price Elasticity, by # of Instruments Selected

Figure displays the estimated own-price elasticity of demand $\hat{\alpha}$ from Equation (5), as a function of the number of instruments selected by the LASSO. We re-solve the first-stage LASSO of Equation 4 along a grid of 101 log-spaced penalty parameters, which are then converted to the corresponding number of selected instruments. Estimation is on a sample of firearm model-years, restricted to firearm models with at least 100 purchases in Massachusetts between 2016–2022. Solid blue line is the point estimate. Dashed red lines are pointwise 95% confidence intervals based on heteroskedasticity-robust asymptotic standard errors. The vertical black line marks the number of instruments selected by the data-driven penalty of our preferred specification.

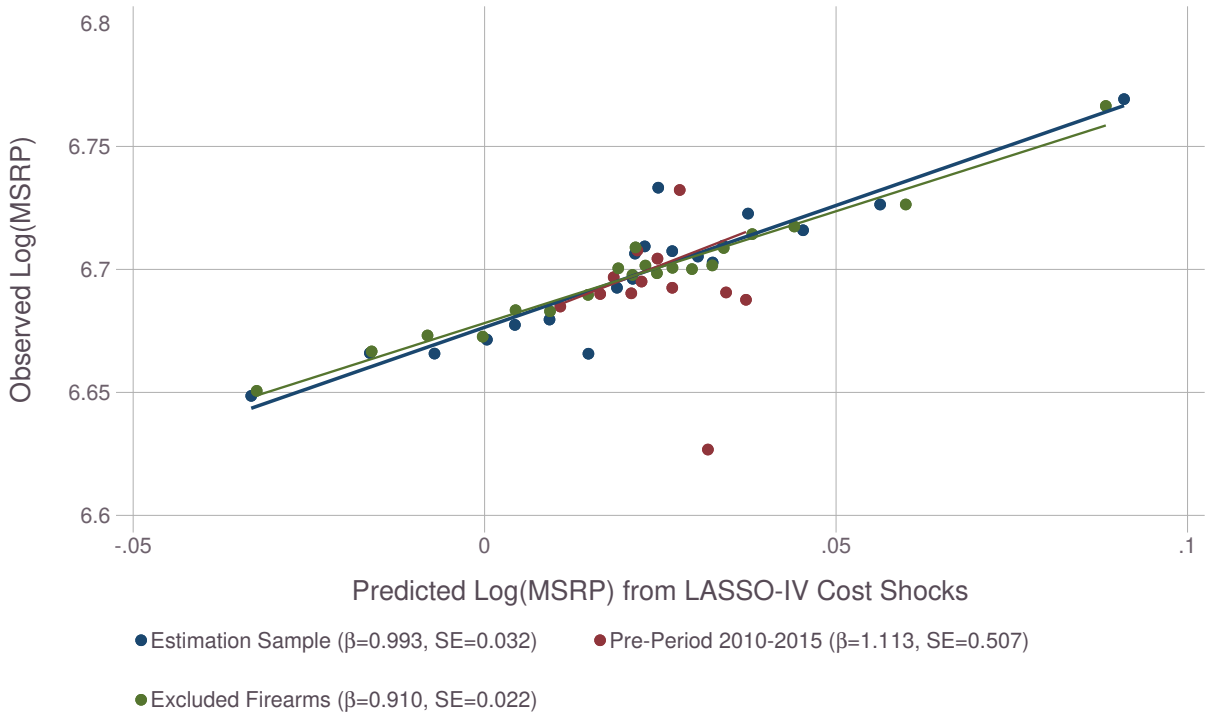


Figure C.4: Predictive Power of Cost Shocks on Alternative Samples

Figure displays binned scatterplot of log MSRP observed for firearm model-years in different datasets against their predicted log MSRP from the input prices of metal commodities under our preferred IVLASSO procedure $\sum_{k \in \mathcal{K}} \hat{\beta}_{m(j),k} \log(\tilde{p}_{k,t})$. All predictions use the same LASSO coefficients $\hat{\beta}_{m,k}$ and metal commodity prices as our preferred approach in Section 4.1. Colors represent different datasets of firearms: (Blue, "Estimation Sample") is the sample used for our preferred approach in Section 4.1, firearm model-years under active production with at least 100 purchases in Massachusetts between 2016–2022; (Red, "Pre-Period") is these same firearm models during 2010-2015, which are years not present in our preferred estimation sample; and (Green, "Excluded Firearms") is a sample from manufacturers in our preferred estimation sample, using firearm models that have fewer than 100 purchases in Massachusetts between 2016–2022, which are models not present in our preferred estimation sample. Points are means within 20 quantile bins of predicted log MSRP, after residualizing both variables on fixed effects for firearm model and weapon class-by-year. Lines are OLS fits within each sample. Reported coefficients and heteroskedasticity-robust standard errors are from separate univariate linear regressions of observed log MSRP on predicted log MSRP, unadjusted for uncertainty in the passthrough parameters $\hat{\beta}_{m(j),k}$ used to construct the prediction.

D Model Estimation Details

D.1 Estimation Routine

We estimate the equilibrium model of the firearms industry in four steps. First, we estimate the mean utilities and parameters governing preference heterogeneity $\Theta = (\rho_0, \rho_1, \alpha, \Pi_\beta, \Sigma_\beta, \Pi_\alpha, \sigma_\alpha, \Pi_\omega, \sigma_\omega)$ via constrained MLE (Goolsbee and Petrin 2004, Train 2009) using the FRB transaction microdata in Massachusetts. Next, we estimate the baseline parameters governing tastes for characteristics β via a projection of δ_j on the time-invariant characteristics X_j. That is, we run the linear regression specified in Equation 7, then use the estimated fixed effects $\hat{\delta}_j$ to estimate β via a GLS regression, with weights proportional to the inverse covariance matrix of fixed effects $\hat{V}$ (Nevo 2000). Next, we estimate the market-specific tastes for firearms $\tau_{s,t}$ across state-years using the adjusted NICS background check data, which captures differences in aggregate tastes for firearms across state-year. Finally, we recover the marginal costs of firearm production $c_{j,t}$ using the first order conditions of the supply side, given the estimated demand parameters.

We collapse the Massachusetts FRB data by gender-zip code demographic cell, denoted d, each year t. Let $\mathcal{T}_{MA}$ denote the markets (state-years) in Massachusetts, $N_{j,d,t}$ the number of individuals purchasing product j in demographic cell d, and $\mathcal{D}_{MA}$ the set of these demographic cells. Let $V_{i,j,s,t} = -\alpha_i \cdot p_{j,t} + \vec{\beta}_i \cdot \vec{X}_j + \delta_{j,s,t}$, the component of utility unrelated to the shock. Conditional on (α_i, β_i), determined by the observed and unobserved heterogeneity, choice probabilities take on the familiar nested logit form:

$$Pr(j|i,s,t) = \frac{\exp(V_{i,j,s,t}/\rho_1)}{\sum_{k \in \mathcal{J}_t, c(k)=c(j)} \exp(v_{i,k,s,t}/\rho_1)} \cdot \frac{\exp(\frac{\rho_1}{\rho_0} IV_{i,c(j),s,t})}{\sum_{c' \in \{h,l\}} \exp(\frac{\rho_1}{\rho_0} IV_{i,c',s,t})} \cdot \frac{\exp(\rho_0 IV_{i,1,s,t})}{1 + \exp(\rho_0 IV_{i,1,s,t})} \tag{D.1}$$

$$= Pr(j|i,s,t,c(j)) \cdot Pr(c(j)|i,s,t,j \neq 0) \cdot Pr(j \neq 0|i,s,t) \tag{D.2}$$

$$IV_{i,c,s,t} = \log\Big(\sum_{j \in \mathcal{J}_t, c(j)=c} \exp(V_{i,j,s,t}/\rho_1)\Big) \tag{D.3}$$

$$IV_{i,1,s,t} = \log\Big(\sum_{c \in \{h,l\}} \exp(\frac{\rho_1}{\rho_0} IV_{i,c,s,t})\Big) \tag{D.4}$$

which can be broken out into 3 distinct choice probabilities: the probability of choosing j conditional on class c $Pr(j|i,s,t,c(j))$, the probability of choosing class c conditional on choosing an inside good $Pr(c(j)|i,s,t,j \neq 0)$, and the probability of choosing an inside good $Pr(j \neq 0|i,s,t)$. $IV_{i,c,s,t}$ and $IV_{i,1,s,t}$ denote the inclusive values of each class and the inside goods, respectively. Conditional on d, consumers in each cell differ only due to the normally

distributed unobserved heterogeneity, so demographic cell market shares are as follows:

$$Pr(j|d,s,t) = \int_{\vec{\nu}} Pr(j|D_i = d, t, \nu_i = \nu, s)\varphi(\vec{\nu})d\vec{\nu} \tag{D.5}$$

where φ represents the standard multivariate normal pdf. As this high-dimensional integral has no closed form, we approximate its value using sparse grid quadrature (Conlon and Gortmaker 2020). We solve the following constrained maximum likelihood problem:

$$\max_{\Theta} \sum_{t\in\mathcal{T}_{MA}} \sum_{d\in\mathcal{D}_{MA}} \sum_{j\in\mathcal{J}_t} N_{j,d,t} \log(Pr(j|d,s,t)) \tag{D.6}$$

$$\text{Subject to: } Pr(j|s,t) = \hat{s}_{j,MA,t} \quad \forall j \in \mathcal{J}_t, t, s = MA \tag{D.7}$$

$$E[\xi_{j,t} \cdot \vec{Z}_{j,t}] = 0 \tag{D.8}$$

$$E[\eta_{c,t} \cdot \tilde{J}_{c,t}] = 0 \quad \forall c. \tag{D.9}$$

This is a standard MLE problem, with three sets of additional constraints which we explain below. Equation D.7 constrains the predicted state-level choice probabilities $Pr(j|s,t)$ to match the empirically observed choice probabilities/market shares $\hat{s}_{j,MA,t}$ in the Massachusetts FRB data (Goolsbee and Petrin 2004). Given Θ, $Pr(j|s,t)$ can be calculated by integrating over the distribution of demographic cells d in Massachusetts.

$$s_{j,MA,t} = Pr(j|t, s = MA) = \frac{\sum_{d\in\mathcal{D}_{MA}} M_{d,t} Pr(j|d,s,t)}{\sum_{d\in\mathcal{D}_{MA}} M_{d,t}} \tag{D.10}$$

where $M_{d,t}$ is the market size (potential gun consumers) belonging to demographic cells d. The empirical analogue to this model-derived probability is the quantity $\hat{s}_{j,MA,t}$, the empirically observed probability of gun j being chosen in market t:

$$\hat{s}_{j,MA,t} = \frac{\sum_{d\in\mathcal{D}_{MA}} N_{j,d,t}}{\sum_{d\in\mathcal{D}_{MA}} M_{d,t}} \tag{D.11}$$

This constraint exactly identifies the mean utilities $\delta_{j,s,t}$ across products in Massachusetts. We implement this constraint via the contraction mapping for nested logit models suggested in Conlon and Gortmaker (2020).

A threat to identification of Θ, and in particular α, is that the national MSRP $p_{j,t}$ may be correlated with the unobserved product-year demand shock $\xi_{j,t}$, as firms strategically set higher prices during heightened demand. For this reason, we also impose during optimization the moment in Equation D.8, that the demand shock is orthogonal to a matrix of instruments. For identifying α, we use the same (residualized) manufacturer specific commodity response

functions $Z_{1,j,t}$ described in Section 4.1[83]. This ensures that α is identified off of price shifts attributable to supply-side cost shocks. It is analogous to the moments used in Berry et al. (1995) to identify price sensitivity in a method of moments demand estimation framework.

Identification of the nesting parameters ρ_1, ρ_0 requires instruments that shift the market shares of products conditional on class $Pr(j|i,s,t,j \in c)$, and weapon class shares conditional on purchase $Pr(j \in c|i,s,t,j \neq 0)$, respectively (Verboven 1996). To identify ρ_1, we include in $\vec{Z}_{j,t}$ a variant of the differentiation instruments proposed in Gandhi and Houde (2019), denoted $Z_{2,j,t}$, that measures the relative impact of cost shocks with respect to rival goods in the same nest:

$$Z_{2,j,t} = \sum_{k \in c(j), f(j) \neq f(k)} (Z_{1,j,t} - Z_{1,k,t})^2 \tag{D.12}$$

This instrument captures the impact of cost shocks on demand for product j, *relative* to its competitors owned by different firms f that are in the same class, and should be negatively correlated with conditional product shares.

Our final constraint (Equation D.9) uses variation in the choice set size to identify the outside option nesting parameter ρ_0. We implement this identification econometrically by enforcing that the time series of preference shocks $\eta_{c,t}$ is orthogonal to time series variation in the (de-meaned) count of firearm models within each class, $\tilde{J}_{c,t} = |\mathcal{J}_{c,t}| - \frac{1}{|T|}\sum_{t'} |\mathcal{J}_{c,t'}|$. This moment condition extends the approach of identifying the nesting parameter based on the entry and exit of goods from the choice set to our setting with multiple levels of nesting (e.g., Miller and Weinberg 2017, Gandhi and Nevo 2021). Choice set size, *ceteris paribus*, should be positively correlated with weapon class shares conditional on purchase.

To interpret the choice set variation that identifies the nesting parameter ρ_0, Panel (a) of Appendix Figure D.1 shows the variation from year-to-year in the size of the class-specific choice set $|\mathcal{J}_{c,t}|$ over our sample. Much of the variation in the choice set size comes from eight merger and acquisition activities occurring across the 58 manufacturers in our sample. For example, in Panel (b), we plot a case study of the change in the number of models produced (according to the Blue Book of Gun Values) by Remington and Marlin Firearms, two manufacturers owned by Remington Outdoor Company. In 2020, Remington Outdoor Company filed for bankruptcy, and the two manufacturers were sold to different companies (Sturm Ruger and Roundhill Group).[84] Collectively, these manufacturers accounted for approximately 5% market share of new guns sold in Massachusetts prior to the bankruptcy in

[83] Residuals are taken by regressing the instrument on product and class-year fixed effects. To increase the power of these instruments, we perform the IVLASSO procedure on all products seen in the BBGV data that share a manufacturer with a product in the demand model. Note also that the composite goods ω_c are excluded from this expectation.

[84] Source: https://www.nytimes.com/2020/07/28/business/remington-bankruptcy-guns.html

2019, dropping to less than 1% by 2021, in part because Roundhill Group delayed production and eventually closed a major firearms plant.[85]

To identify the parameters governing demographics-driven heterogeneous tastes for gun characteristics $\Pi_\beta, \Pi_\alpha, \Pi_\omega$, we use the differential buying patterns within a product across consumers in different gender-zip code cells. We identify the scale of the unobserved heterogeneity $\Sigma_\beta, \sigma_\alpha, \sigma_\omega$ via the entry and exit of products with differing characteristics.

Finally, to recover $\tau_{s,t}$ in each state-year, we match the demand model's prediction to the aggregate number of firearms purchased, according to the adjusted NICS checks from Section B.3. We solve for the following root in each state-year:

$$\tau_{s,t} : \hat{s}_{1,s,t} = \int_d \int_\nu \frac{\exp(\tau_{s,t} + \rho_0 IV_{i,1,s,t})}{1 + \exp(\tau_{s,t} + \rho_0 IV_{i,1,s,t})} \varphi(\nu) dF(d|s,t) \tag{D.13}$$

where i indexes the demographics d and random heterogeneity ν governing individual tastes, and $IV_{i,1,s,t}$ is the corresponding inclusive value of inside goods for i.

Given our demand estimates, we recover the marginal costs of production for each firearm model using the firm's first order conditions for profit maximization. Given the profit objective in Equation 8, the first order condition for firm f's profit maximization is as follows:

$$\frac{\partial \Pi_f}{\partial p_j} = (1 - v_j) q_{j,t} + \sum_{k \in \mathcal{J}_{f,t}} ((1 - v_k) p_{k,t} - c_{k,t}) \frac{\partial q_{k,t}}{\partial p_j} = 0 \tag{D.14}$$

Re-arranging this first order condition, we obtain an expression for the marginal cost of firearms production:

$$\vec{c}_f = (1 - \vec{v}_f) \odot \vec{p}_f + \Delta_f'^{-1} ((1 - \vec{v}_f) \odot \vec{q}_f) \tag{D.15}$$

where Δ_f is the $|\mathcal{J}_f| \times |\mathcal{J}_f|$ matrix of cross-price derivatives (i.e. $\Delta_{j,k} = \partial q_j / \partial p_k$). Given estimates of demand, we can recover the marginal cost directly from the firm's first-order condition.[86]

D.2 Model Validation

In this section, we provide details on out-of-sample moments we compare our model estimates to for validation. In Figure D.7, we assess the fit of our demand model, using out-of-sample moments from four distinct data sources. In Panel (a), we use additional information from

[85] Source: https://www.wktv.com/news/top-stories/remarms-ilion-operation-to-close-march-2024/article_55edb272-8fbd-11ee-9b84-1b80825a8ee7.html

[86] This procedure yields 2% of product-years having negative marginal cost. These products constitute relatively niche products in the demand system: they account for 0.1% of all purchases of new firearms in our sample period.

NICS on the share of hand versus long gun background checks in each state, allowing us to compare the fraction of handgun purchases observed in each state against the fraction predicted by our model, conditional on any purchase.[87] This is an out of sample test, since we only use total NICS checks in other states to estimate our demand parameters, not checks broken down by firearm class. Therefore, variation in the predicted share of handgun purchases is driven solely by the demographic differences across states, capturing the extent to which our estimates of demographic preferences from Massachusetts can explain handgun purchase rates in other states. Our model is unable to exactly fit the data, as the model's predictions are less dispersed than the observed distribution of handgun purchase rates. This may be due to differences in laws concerning handgun purchases across states, which are not accounted for in our model, or differences in class reporting across state lines to NICS, as states are given significant leeway in how they report these sub-items (Smucker et al. 2022). At the same time, the model predicts patterns in the right direction: the correlation between our predicted handgun purchase rate and the observed handgun background check rate is 0.40, and apart from Hawaii, there are no notable outliers.

In Panel (b), we compare the predicted total quantity of new gun purchases attributable to each firm in our model, to the total number of new guns manufactured by each firm per year, according to the ATF. As with the background check data, we would not expect a one-to-one relationship between guns sold and guns produced, since it may take time for guns to travel from the manufacturer to the end user. Nonetheless, we estimate a strong correlation between our model's predictions and the data: the correlation in the log quantities of the model predictions for each firm and the ATF data is 0.75. Importantly, there does not appear to be a systemic under or over prediction of the number of guns produced by each firm.

In Panel (c), we plot the implied excise tax revenue over time for newly produced guns $\sum_{j \in \mathcal{J}_t^N} p_j q_j v_j$ from our structural model, along with the actual excise tax revenue collected by the U.S. government by quarter.[88] We observe very similar trends, though the levels of revenue from our model are somewhat higher (640M$ in FAET, versus 800M$ implied annually by our model), particularly in the earlier part of the sample before 2020.

Since we rely on cross-neighborhood heterogeneity in demographics from Massachusetts to identify demographic preferences β_i, α_i, our extrapolation relies on this heterogeneity across zip codes in Massachusetts to reflect national patterns. Since Massachusetts is one of

[87] We exclude Nebraska, Iowa, and North Carolina from the figure, as these states are partial permit states that do not require background checks for handguns if the purchaser has a permit. See Section B.3 for details.

[88] Data collected from https://www.ttb.gov/system/files?file=images/foia/xls/Quarterly-breakdown-of-FAET-collections.xlsx. Note that we lag FAET revenue by one quarter to reflect the revenue for corresponding sales.

the most left-leaning states, we worry that estimates of preferences along political dimensions may not be reflective of the broader United States.[89] To test this, we compare in Panel (d) estimates of τ, the state-year taste shifter, to the % in each state voting conservative in the 2016 presidential election, to validate our estimates of demographic heterogeneity. Recall τ is the residual preference for firearms *after* accounting for demographics. If purchasing patterns in conservative neighborhoods in Massachusetts are informative about patterns in conservative neighborhoods nationally, we would expect that this shock is uncorrelated with the political leaning of the state. We find that there is very little correlation between τ and conservative voting patterns in the state ($\rho = -0.01$), suggesting our demographic preferences in Massachusetts along this dimension are representative of the greater United States population.

This suggests our model is able to accurately capture national patterns in firearm demand, despite using only data on aggregate quantities from other states to identify our demand model.

[89] Along other demographics in our model, the distribution in Massachusetts is quite similar.

Consumer Characteristic:	Baseline	Female	Frac White	Frac Conservative	Log(Median Inc)	Log(Density)	Nesting Parameter (ρ)	Random Coef. (σ)
Handgun (vs No Gun)	-0.719***	-1.182***	0.965***	1.179***	-0.433***	-0.113***	0.922	
	(0.005)	(0.044)	(0.043)	(0.041)	(0.012)	(0.002)	(0.094)	
Long gun (vs Handgun)	-0.270***	-0.717***	0.519***	0.139***	-0.028***	-0.011***	0.268***	
	(0.006)	(0.087)	(0.071)	(0.030)	(0.010)	(0.002)	(0.010)	
Shotgun (vs Rifle)	-0.073***	0.090***	-0.108***	-0.165***	0.137***	-0.027***		0.074
	(0.007)	(0.017)	(0.037)	(0.051)	(0.015)	(0.003)		(0.559)
Used Gun (vs New for $0)	1.251***	-0.798***	0.571***	0.185***	0.037***	0.035***		0.729***
	(0.035)	(0.032)	(0.040)	(0.039)	(0.012)	(0.002)		(0.051)
Caliber (In.)	0.241***	-0.077***	0.178**	0.502***	-0.232***	0.081***		0.042
	(0.014)	(0.027)	(0.071)	(0.093)	(0.028)	(0.006)		(1.961)
Barrel Length (Ft.)	0.096***	-0.326***	0.201***	0.092***	-0.064***	0.002**		0.027
	(0.004)	(0.014)	(0.020)	(0.020)	(0.006)	(0.001)		(0.271)
High Capacity Weapon	-0.386***	-0.209***	-0.081***	0.081***	0.057***	0.011***		0.789***
	(0.002)	(0.010)	(0.016)	(0.021)	(0.006)	(0.001)		(0.033)
Price Coef. (Multiplier)	0.053	0.567***	-0.077***	-0.267***	-0.279***	-0.012***		0.458***
	(0.048)	(0.015)	(0.018)	(0.025)	(0.009)	(0.001)		(0.015)

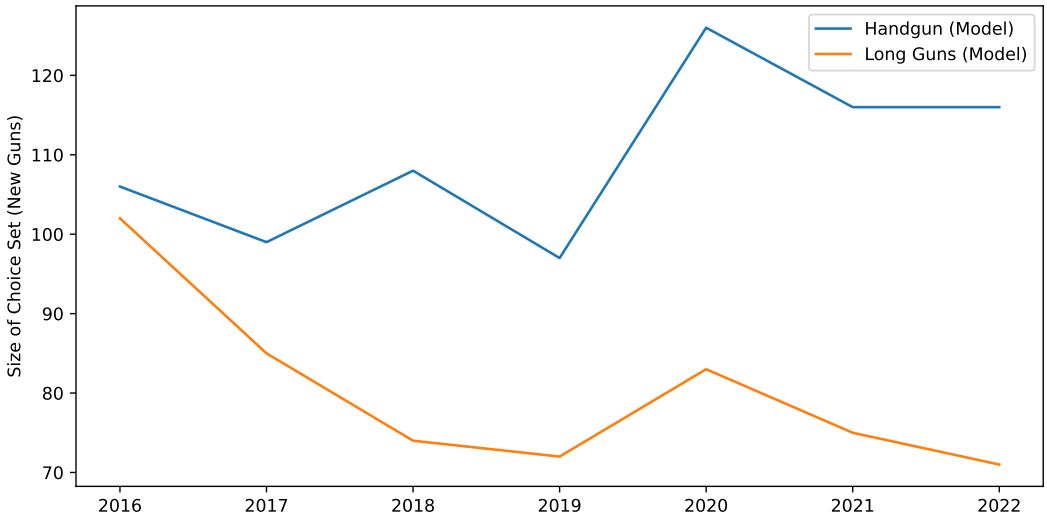
Table displays estimates of the non-linear parameters of our demand model from Section 5, using the procedure detailed in Appendix D. Nesting parameter in the top row is ρ_0. Nesting parameter in the next row is ρ_1. Standard errors in parentheses and calculated using corrected formula for constrained maximum likelihood described in Moore et al. (2008). Stars indicate statistical significance of p-values. For all values besides the nesting parameters $\theta \in \Theta \setminus \{\rho_0, \rho_1\}$, stars correspond to significance of the test $Pr(\theta) = 0$. For the nesting parameter ρ_0 between inside and outside options, the test is $Pr(\rho_0 < 1)$. For the nesting parameter ρ_1 between inside options within a firearm classes, the test is $Pr(\rho_1 < \rho_0)$. These significance tests on the nesting parameters are performed using a one-sided Z-test, treating the upper bound of the test as a known number. * denotes $p < .1$, ** denotes $p < .05$, *** denotes $p < .01$,

Table D.1: Demand Parameters

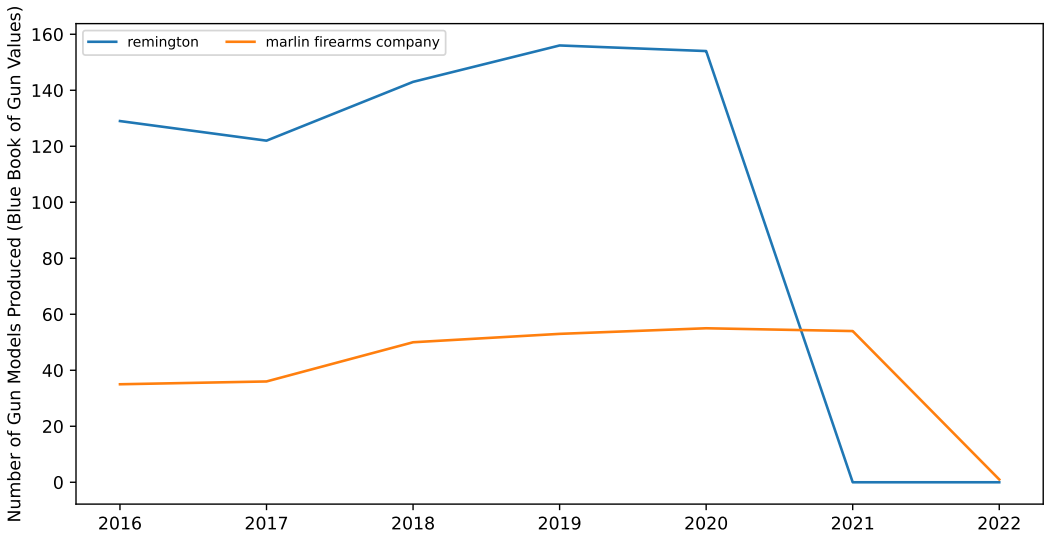
	Gender	10% Increase in Neighborhood Characteristic			
Demographic	Female	Log(Median Inc)	Frac White	Log(Density)	Frac Conservative
Handgun (vs No Gun)	-413.93	-71.81	78.75	-11.64	76.17
Shotgun (vs Rifle)	113.18	10.89	-10.24	-2.50	-16.44
Long gun (vs Handgun)	-372.79	-16.28	43.38	-1.59	-0.43
Caliber (In.)	-228.12	-13.36	18.32	7.74	53.02
Barrel Length (Ft.)	-285.36	-6.32	18.12	0.17	7.82
High Capacity Weapon	24.36	-5.33	-10.30	0.57	-2.68
Used Gun (vs New for $0)	-1,253.23	29.27	59.06	4.34	41.62

Table displays estimates of the average change in Massachusetts consumers' willingness-to-pay for a firearm characteristic x ($WTP_{i,x} = \beta_{i,x}/\alpha_i$), with respect to their demographics D_i, ($\partial WTP_{i,x}/\partial D_i$). For neighborhood characteristics, we scale this derivative by 0.1 to represent a .1 log point (roughly 10%) increase in the demographic of interest for logged characteristics (density and median income), and a 10 percentage point increase for characteristics in levels (fraction white and conservative). We compute this value by evaluating the derivative within each demographic-zip code cell, and taking a weighted average across cells, with weights proportional to each cell's market size in Massachusetts. For each cell, we compute the average across i weighting by the market size of each respective demographic cell. All values in 2022 dollars.

Table D.2: Change in Willingness-to-pay for Gun Characteristics, by Demographic (in $)



(a) Choice sets $\mathcal{J}_{c,t}$ by Weapon Class and Year



(b) Case Example: Remington's Exit from Firearm Production

Figure D.1: Choice Sets in Demand Model over Time

Figure displays year-to-year variation in the choice sets of consumers in our demand model. Panel (a) plots the number of distinct firearm models entering the actively-produced ("new") gun choice set $\mathcal{J}_t^N$ each year, separately for handguns and long guns. Panel (b) plots the number of distinct gun models listed in active production in the Blue Book of Gun Values each year for Remington and Marlin, the two firearm brands owned by Remington Outdoor Company at the time of it bankruptcy filing in 2020.

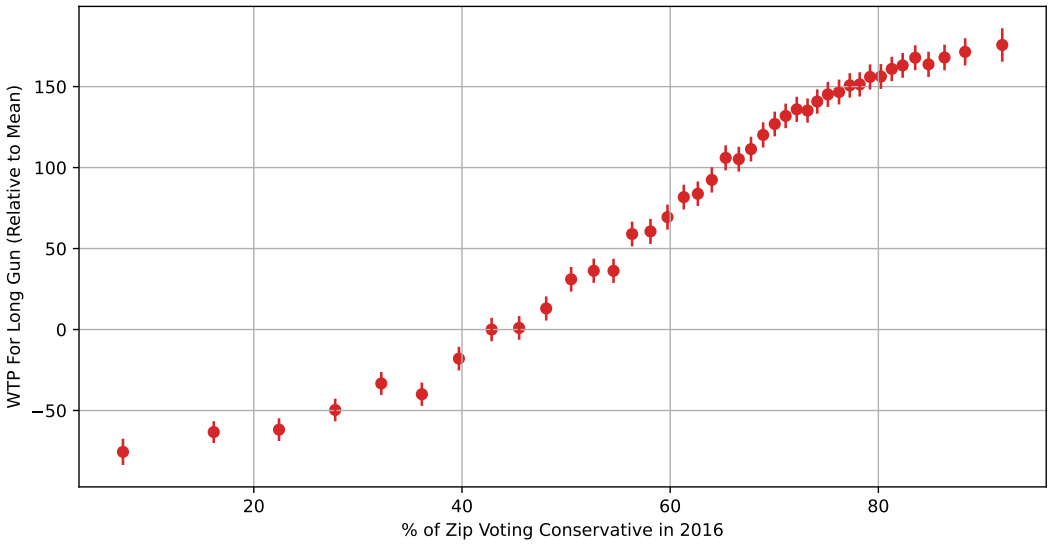


Figure D.2: Willingness to Pay for Long Guns Across Political Preferences

Figure displays a binned scatterplot of each zip code's 2016 presidential vote going to conservative candidates and its consumers' average willingness-to-pay to purchase a long gun instead of a handgun. For each demographic cell, we its consumers' average WTP for the long gun characteristic, $\beta_{i,l}/\alpha_i$, where both terms include the demographic interactions, and we integrate over the distribution of random coefficients. WTP is expressed in 2022 dollars and normalized to zero at the national average conservative vote share, with unweighted average across zip codes of 44%. The procedure of Cattaneo et al. (2019) is used to select the number and location of points and to form 95% confidence intervals.

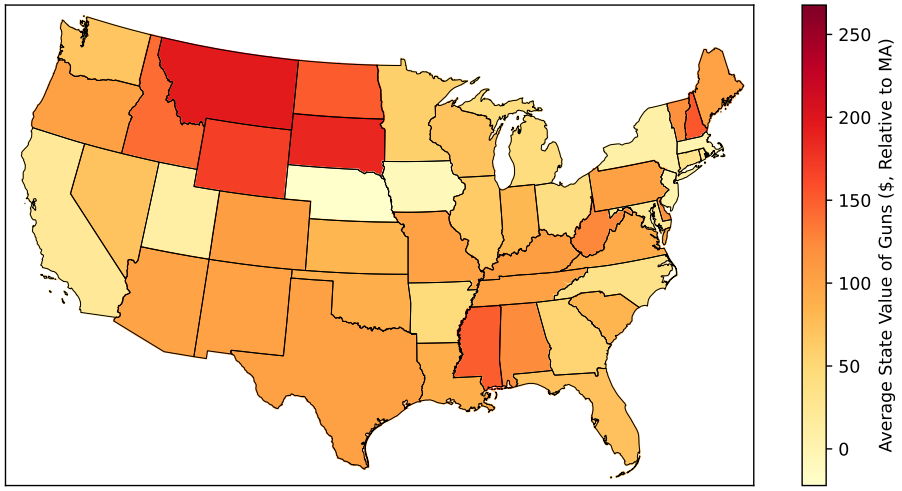
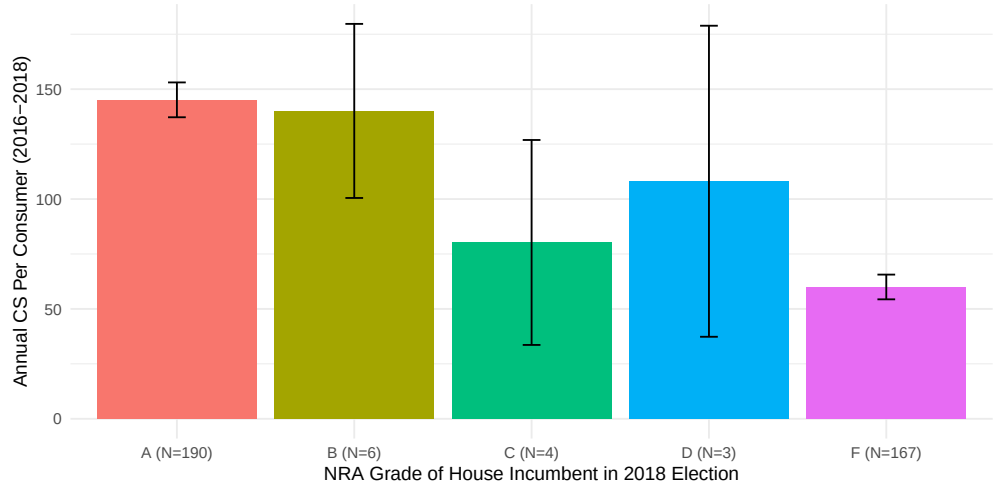
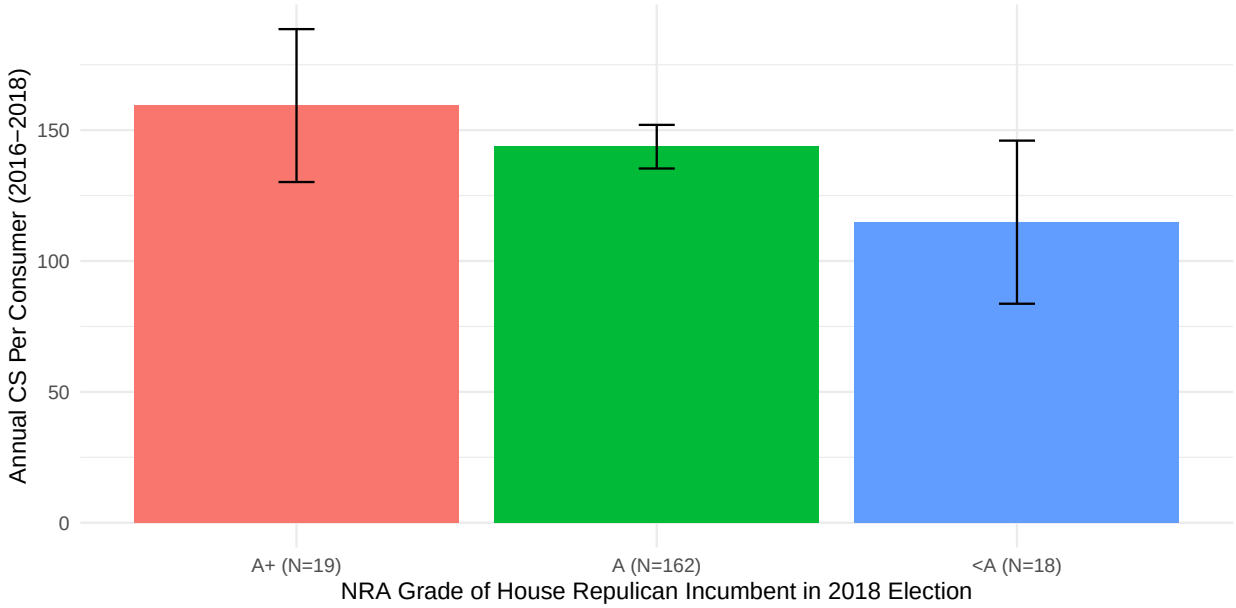


Figure D.3: Average State-year Taste Shifters for Firearms by State, in $

Figure shows the value of $E_s[CS_{i,t}] - E_s[CS_{i,t}|\tau = 0]$, where the expectation is computed across consumers, separately for each state-year, and then averaged overs years t within each state s. This difference expresses the average consumer's surplus from the firearm market in each state, relative to an otherwise identical state (e.g., product assortment, demographic distribution, changes over time) with Massachusetts' vertical firearm preference shock $\tau_{MA,t} = 0$. Consumer surplus is calculated analogously to Equation (F.3).



(a) NRA Grades for Incumbent Congress Members



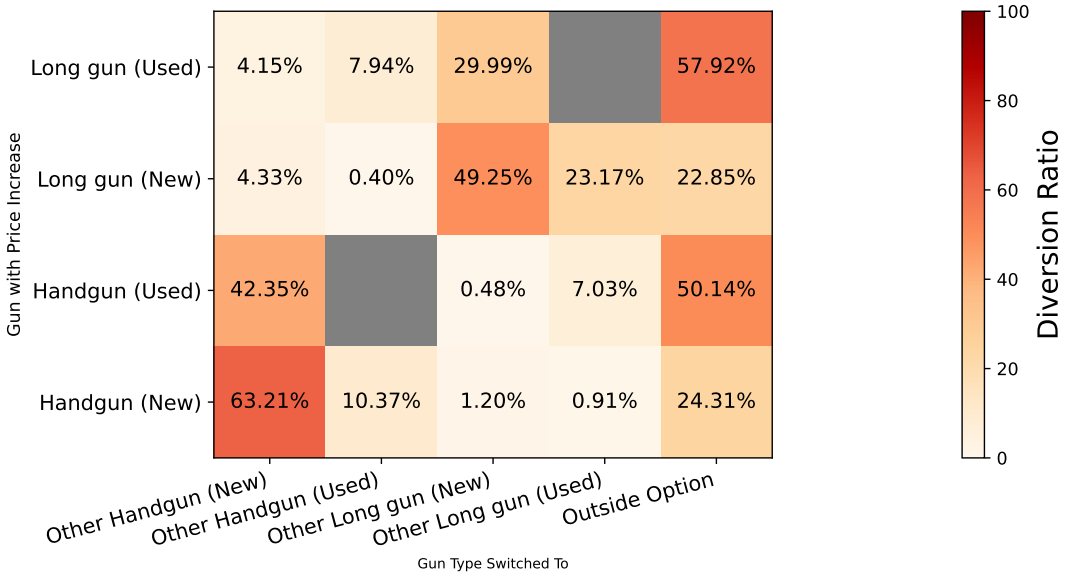
(b) Consumer Surplus and NRA Congressional Grades, Among Republican Seats

Figure D.4: Consumer Surplus and NRA Grades

Figure shows the relationship between NRA grades of congressional representatives and a measure of consumer surplus from the firearm industry among their constituents, across U.S. House Congressional Districts. Panel A displays all House districts' average in-sample consumer surplus per consumer between 2016–2018, across letter grades given by the NRA to each district's House Representative during the 2017-2018 session. Panel A displays Republican-represented House districts' average in-sample consumer surplus per consumer between 2016–2018, across letter grades given by the NRA to each district's House Representatives during the 2017-2018 session. Both panels include only incumbent representatives running for re-election that have been assigned a letter grade by the NRA. Consumer surplus is calculated analogously to Equation (F.3).



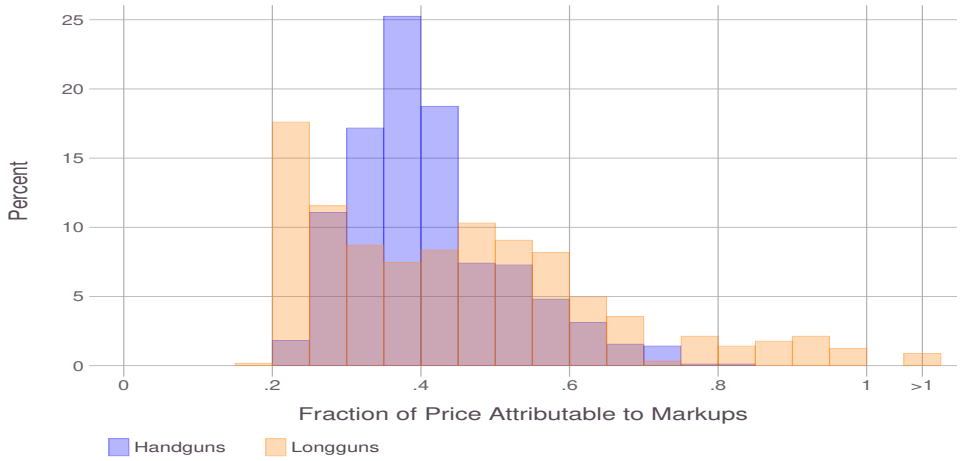
(a) Distribution of National Own-Price Elasticities, by Gun Class



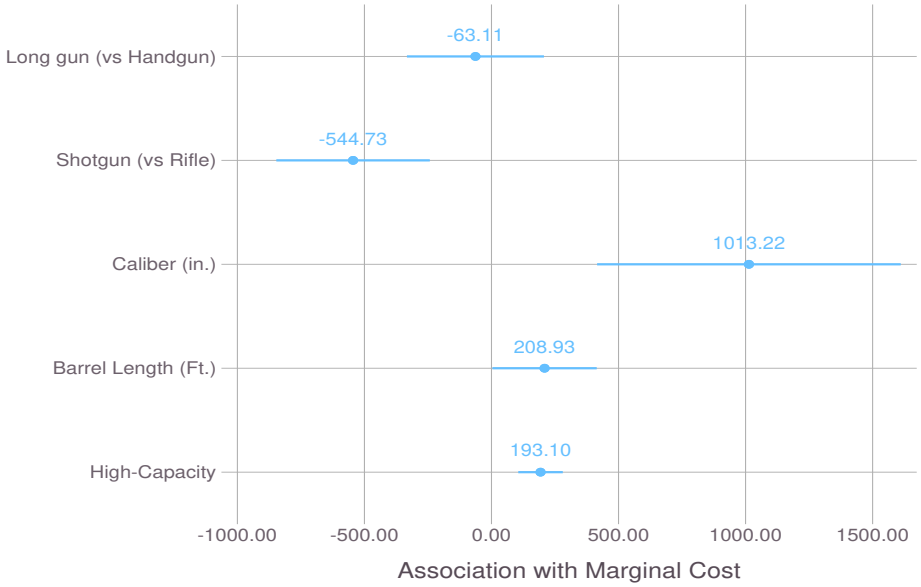
(b) National Substitution Patterns (Diversion Ratios) across Gun Types

Figure D.5: Own and Cross Firearm Substitution Patterns

Panel (a) displays the distribution of national own-price elasticities $\partial \log(q_{j,t})/\log(p_{j,t})$ across new gun model-years under active production, by weapon class. Panel (b) displays the average diversion ratio across groups of firearms, with diversion ratio for group G defined as $-(\sum_{k\in G,k\neq j}\partial q_{k,t}/\partial p_{j,t})/(\partial q_{j,t}/\partial p_{j,t})$. Diversion ratios are weighted by the national sales quantity of each alternative j within group G.s



(a) Distribution of Lerner Indices, by Gun Class



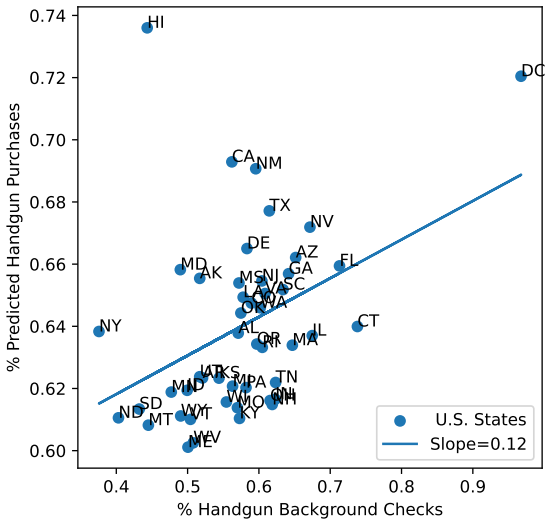
(b) Correlates of Marginal Costs with Gun Characteristics

Figure D.6: Marginal cost estimates

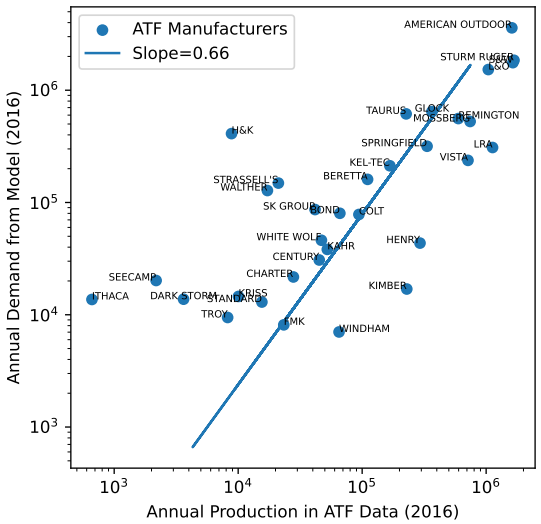
Panel (a) displays the distribution of estimated Lerner indices across firearm model-years under active production, by weapon class. Lerner index is defined as the share of revenue from a firearm sale accruing as profit to the firm $(p_{j,t}(1-v_j)-c_{j,t})/p_{jt}$. Panel (b) displays the estimated correlates of marginal cost from the following two-step regression:

$$c_{j,t} = c_j + \upsilon_t + \epsilon^c_{j,t}$$
$$c_j = \beta X_j + \epsilon^c_j$$

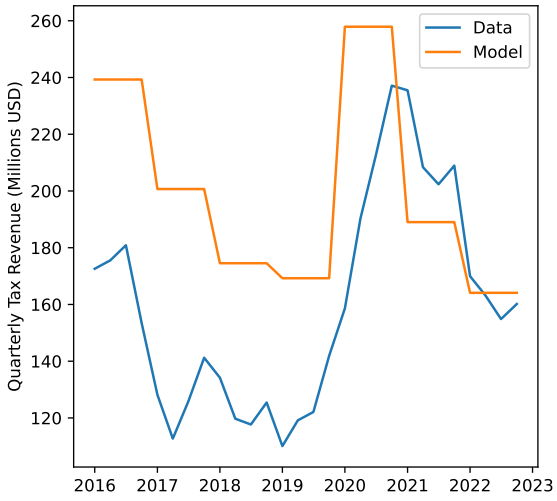
Where $c_{j,t}$ is estimated from the first order conditions of firm profit maximization, and the second regression is estimated via GLS with weights proportional to the inverse variance of the estimated fixed effects (Nevo 2000).



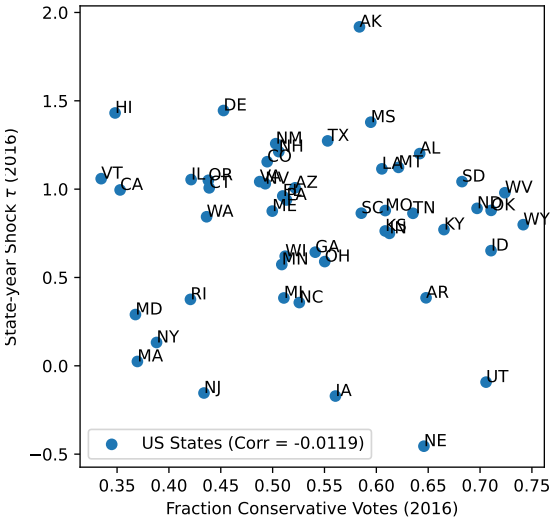
(a) % Handgun Sales by State (Model vs NICS Data)



(b) 2016 Sales by Domestic Firms (Model vs ATF Production Data)



(c) Excise Tax Revenue by Quarter (Model vs FAET Filings)



(d) % Conservative vs State-year Shock (2016)

Figure D.7: Out-of-Sample Moments of National Demand vs Model Predictions

Figure displays model predictions of out-of-sample quantities, described in Section D.2. Panel (a) displays the in-sample average share of each state's adjusted NICS background checks that are attributable to handgun purchases, relative to the in-sample average of our model's predicted share of all gun purchases that are handguns during the sample. Panel (b) displays the in-sample average predicted quantity of guns purchased from each manufacturer, compared to the in-sample average number of guns produced by these manufacturers and sold into the consumer firearms market, according to the ATF. Both panels (a) and (b) include a best-fit line via OLS. Panel (c) displays the predicted quarterly tax revenue from the model $\frac{1}{4}\sum_{j\in\mathcal{J}_t^N} q_{j,t}\cdot p_{j,t}\cdot v_{c(j)}$ (we divide the predicted annual revenue by 4), compared to the actual revenue collected by the FAET. Panel (d) displays the fraction of votes for conservative candidates in the 2016 Presidential Election, relative to the state-level demand shock estimated from our model in 2016 $\hat{\tau}_{s,2016}$.

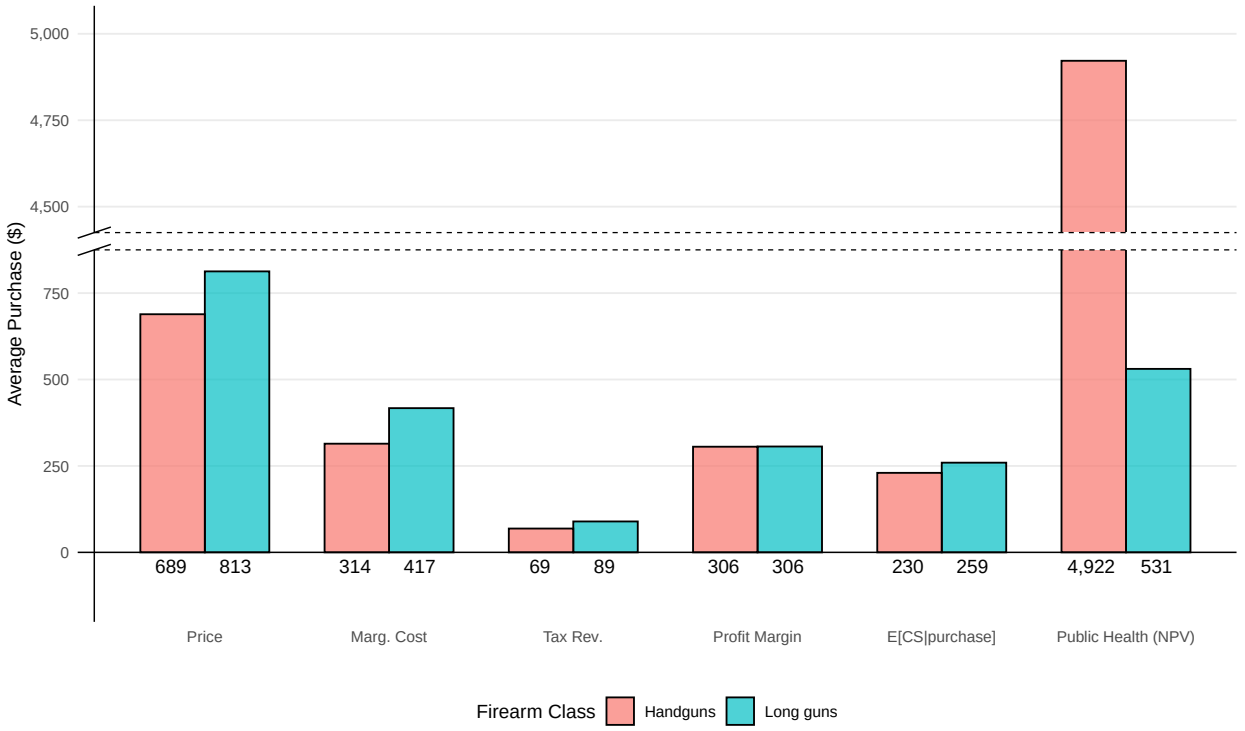


Figure D.8: Distribution of Surplus from Firearm Purchase, Quantity-Weighted

Figure displays distribution of surplus in dollars from the average marginal firearm purchase in each class, weighted by model-predicted firearm purchase quantities in the U.S. between 2016–2022. Price is the firearm's MSRP $p_{j,t}$. Tax Revenue is the revenue to the federal government generated by the firearm's sale under status quo regulation $v_j p_{j,t}$. Marginal cost is our estimate of the marginal production cost $c_{j,t}$. Profit margin is our estimate of price, less taxes and marginal cost $\mu_{j,t} = p_{j,t}(1 - v_j) - c_{j,t}$. By construction, the sum of tax revenue, marginal cost, and profit margin equals the price. Consumer surplus measures the average value to consumers from purchasing the firearm, conditional on that firearm being the best alternative in their choice set. Public health is equal to the net present value of expected lifetime gun homicides caused by a firearm purchase, converted to dollars via a value of a statistical life. Vertical axis suppresses \$1,500–\$4,000 for readability.

E Details on Model of the Firearm Stock and Homicides

E.1 Model Structure and Microfoundation

We microfound in this section the transmission of firearms stock to homicides, taking the magnitude of the stock $Q_{c,s,t}$ as given. We assume that additional firearms generate (fatal or non-fatal) assaults, and these assaults translate to homicides with some probability.

We assume that the rth firearm in the stock $r \in \{1, 2, \ldots, Q_{c,s,t}\}$ stochastically generates firearm assaults a_r attributable to class c according to the following distribution

$$a_r \sim \text{Pois}\left(\left(r^\kappa - (r-1)^\kappa\right)\chi_c^a \zeta_{st}\right),$$

purely independently. Here, $r^\kappa - (r-1)^\kappa$ is the marginal risk of a class-c firearm assault added from the rth firearm in the stock, relative to the $r = 1$th. This marginal risk is strictly positive and decreasing, capturing the idea that any firearm added to the stock generates additional assaults with some probability, but less so when the pre-existing stock is already large. The parameter χ_c^a further scales this marginal risk across different classes (e.g., handguns are easier to carry and conceal than long guns). The parameter ζ_{st} is a state-year shock reflecting the risk that any firearm in the stock—marginal or inframarginal—is used to perpetrate an assault (e.g., due to a change in policing or criminal justice policy).

Then using the additive property of the Poisson, the count of assaults in each class-state-year follows the distribution

$$a_{c,s,t} \sim \text{Pois}\left(\sum_{r=1}^{Q_{cst}} \left(r^\kappa - (r-1)^\kappa\right)\chi_c^a \zeta_{st}\right) = \text{Pois}\left(Q_{cst}^\kappa \chi_c^a \zeta_{s,t}\right).$$

Following the literature (Braga et al. 2021), we assume that the risk a shooting turns fatal is unrelated to individual intent or offender characteristics (but may be related to firearm characteristics). Thus, the occurrence of fatalities by class-state-year has the compound distribution

$$d_{c,s,t} \sim \text{Binom}\left(a_{cst}, \chi_c^d\right),$$

where a_{cst} is the number of firearm assaults each class-state-year, and χ_c^d is the probability that each one results in a fatality. The fatality probability is class-specific as there are documented technological differences between handguns and long guns associated with their

lethality (Libby and Corzine 2007).

Using well-known properties of the Poisson-Binomial compound distribution, the count of total fatalities is distributed as

$$d_{s,t} = d_{h,s,t} + d_{l,s,t} \sim \text{Pois}\Big(\big(Q^{\kappa}_{hst}\chi_h + Q^{\kappa}_{lst}\chi_l\big)\zeta_{s,t}\Big)$$

$$(d_h, d_l)_{s,t} | d_{s,t} \sim \text{Mult}\left(d_{s,t}, \left(\frac{\frac{\chi_h}{\chi_l}Q^{\kappa}_{hst}}{\frac{\chi_h}{\chi_l}Q^{\kappa}_{hst} + Q^{\kappa}_{lst}}, \frac{Q^{\kappa}_{lst}}{\frac{\chi_h}{\chi_l}Q^{\kappa}_{hst} + Q^{\kappa}_{lst}}\right)\right),$$

where $\chi_c = \chi^a_c \chi^d_c$. Thus, the relative lethality "efficiency" χ_h/χ_l reflects both differences in the ways that handguns and long guns are used in the assault generation process (e.g., ease of carrying) and the technological characteristics that make assaults with handguns differentially lethal from those with long guns.

E.2 Moments

E.2.1 Initial Conditions

Our calibration of the initial firearm stock utilizes the following decomposition:

$$Q_{c,s,2015} = \text{HH}_{s,2015} \times \frac{\text{Adults with gun in HH}_{s,2015}}{\text{Adults}_{s,2015}} \times \frac{\text{Gun owners}_{2015}}{\text{HH with gun}_{2015}} \times \frac{\text{Guns}_{c,r,2015}}{\text{Gun Owner}_{r,2015}}.$$

There are direct measurements of the first two terms—the number of households per state and the share of adults living in a household with a firearm—in publications from the ACS and Schell et al. (2020), respectively.

We measure the final two terms using microdata from the 2015 National Firearm Survey (Azrael et al. 2017). To calculate the number of firearm owners per household with at least one firearm, we compute the average number of firearm owners per household, as reported by survey respondents who personally owned a firearm. We use these same respondents to compute the average count of class c firearms per owner. To avoid survey noise in states with few observations, we compute the guns of class c per gun owner at the census region level r.[90]

[90] In practice, the average guns of each class per owner do not vary much by region, so this assumption is relatively unimportant.

E.2.2 Scrap rate φ

Our assumed law of motion of the firearm stock by state-class-year aggregates to the moment in Section 6.2, as

$$\begin{aligned} Q_t &= \sum_s \sum_c (1-\varphi) Q_{c,s,t-1} + q^N_{c,s,t} \\ &= (1-\varphi) \sum_s \sum_c Q_{c,s,t-1} + \sum_s \sum_c q^N_{c,s,t} \\ &= (1-\varphi) Q_{t-1} + q^N_t. \end{aligned}$$

Since the flow of new firearms q^N_t is observed each year, chaining this equality between any pair of measurements of the firearm stock Q_t, $Q_{t'}$ with $t \neq t'$, provides a unique mapping to the scrap rate φ. For calibration, we choose 1994 and 2015 as our measurement years, which represent the longest time series with data on the stock available that predates our sample period.

E.2.3 Elasticity κ

We use two separate measurements of the elasticity of firearm homicides in a state-year with respect to the contemporaneous firearm stock, coming from Duggan (2001)[91] and Cook and Ludwig (2006).[92] We combine these two estimates into a single calibration target by taking their simple average (0.294), and under the assumption of independence, estimate the standard error of this average.

We set κ equal to this target. Because κ governs a class-specific elasticity in our model, it is not immediately obvious that this is equal to the overall elasticity of homicides $d_{s,t}$ with respect to the overall gun stock $Q_{s,t} = Q_{h,s,t} + Q_{l,s,t}$. We show below that this is indeed the case under certain conditions. Let $f_{c,s,t} = Q_{c,s,t}/Q_{s,t}$ denote the fraction of the stock belonging to class c. If we evaluate the overall elasticity under our model, *holding* f fixed,

[91] Table 6, Column 1, $\hat{\kappa} = 0.316$, $\hat{SE}(\hat{\kappa}) = 0.111$.
[92] Table 3, Column 2, $\hat{\kappa} = 0.272$, $\hat{SE}(\hat{\kappa}) = 0.110$.

we get the following equality:

$$\begin{aligned}
\frac{\partial \log\left(E[d_{st}]\right)}{\partial \log\left(Q_{st}\right)}|_{f_{c,s,t}} &= \frac{\partial}{\partial \log\left(Q_{st}\right)} \log\left(\sum_c Q_{cst}^{\kappa} \chi_c \zeta_{st}\right) \\
&= \frac{\partial}{\partial \log\left(Q_{st}\right)} \log\left(\sum_c \left(f_{c,s,t} \exp\left(\log\left(Q_{st}\right)\right)\right)^{\kappa} \chi_c \zeta_{st}\right) \\
&= \frac{1}{\sum_c \left(f_{c,s,t} Q_{st}\right)^{\kappa} \chi_c \zeta_{st}} \sum_c \chi_c \zeta_{st} f_{c,s,t}^{\kappa} \frac{\partial}{\partial \log\left(Q_{st}\right)} \left(\exp\left(\log\left(Q_{st}\right)\right)\right)^{\kappa} \\
&= \frac{1}{\sum_c \left(f_{c,s,t} Q_{st}\right)^{\kappa} \chi_c \zeta_{st}} \sum_c \chi_c \zeta_{st} f_{c,s,t}^{\kappa} \kappa Q_{s,t}^{\kappa} \\
&= \kappa \\
&\overset{\text{calibrate}}{=} 0.294
\end{aligned}$$

where the third equality uses the assumption that the composition of the firearm stock Q_{cst}/Q_{st} is constant across percentage changes in the stock $\partial \log Q_{st}$ as measured in these papers. This is an assumption that would be violated if demand for firearms of different classes were on systematically different trends across states, though both Duggan (2001) and Cook and Ludwig (2006) provide evidence against this possibility.

E.2.4 Relative lethality χ_h/χ_l

We use several methods for calculating the handgun share of firearm homicides $E[d_{c,s,t}/d_{s,t}]$. All our calculations exclude incidents in which the firearm class was not known, and we assume that such data are missing at random.[93] Our preferred calculation takes observations of the share of firearm homicides in which a handgun was used by state-year from the National Violent Death Reporting System, and aggregates across state years using weights based on the count of homicides in the NCHS vital statistics. We calculate analogous ratios using the FBI's Supplemental Homicide Records and using three different weighting schemes for each average: equal weighting by state-year, weights proportional to firearm homicides within the data source in which the firearm class is known, or weights proportional to total firearm homicides in the NCHS vital statistics. Across these six different measurement approaches, the calibration target ranged from 88.9% to 90.8%, and all produced extremely similar parameters of relative lethality. The firearm class is observed in approximately 56 percent of firearm homicides in the NVDRS records and 48 percent of firearm homicides in the SHR

[93] This assumption allows us to (i) impute the class for firearm fatalities within a state-year in which the class is not observed and (ii) drop state-year observations in which the number of firearm homicides and/or their distribution across class is not recorded.

records, with low correlation in missingness across the two data sources.

E.3 Calibration and Inference

E.3.1 Calibration

We calibrate the model through a series of steps. First, we solve for the pre-existing firearm stock $Q_{c,s,2015}$ and the scrap rate φ, neither of which depends on the other. We then use these estimates, combined with predicted flows from our model of the firearms industry to predict firearm stocks in class-state-years $Q_{c,s,t}$ from 2016–2022.

Turning to the homicide process, we calibrate the elasticity κ as above. Conditional on the elasticity and our measurement of the firearm stocks by class-state-years $Q_{c,s,t}$, the final moment in Section 6.2 is a single equation in a single unknown, which we solve numerically for the relative lethality χ_h/χ_l. We normalize the lethality of long guns to be $\chi_l = 1$.

We use these terms to form predictions of the state-year shocks in our model. The above steps are sufficient to predict firearm homicides under our model, contaminated by two sources of multiplicative bias

$$d_{st} = \overbrace{\frac{d_{st}}{E[d_{st}]}}^{\text{Poisson shocks}} \overbrace{\zeta_{st}}^{\text{Structural shocks}} \left(Q_{l,s,t}^{\kappa} + \chi_h Q_{h,s,t}^{\kappa}\right).$$

The first term on the right-hand side is idiosyncratic Poisson variability about its mean, which has expectation 1 by construction. The second term is the structural shock, which may have any scale or correlation structure in comparison to the firearm stock Q_{cst}. Thus, since d_{st} is an unbiased estimate of $E[d_{st}]$, an unbiased estimate of the structural state-year shock is

$$\widehat{\zeta_{st}} = \frac{d_{st}}{Q_{l,s,t}^{\kappa} + \chi_h Q_{h,s,t}^{\kappa}}.$$

E.3.2 Inference

We perform inference for the scrap rate $\hat{\varphi}$ by adjusting for sampling uncertainty in the surveys of adults used to construct the measures of the firearm stock in 1994 and 2015. We accessed the underlying survey data from Cook and Ludwig (1996) and Azrael et al. (2017), from which we compute standard errors for their reported estimates of Q_{1994} and Q_{2015}, respectively. We then bootstrap our procedure for calibrating the scrap rate $\hat{\varphi}$, drawing replicates Q_{1994}^{b} and Q_{2015}^{b}. In forming the replicates, we draw from a bivariate normal distribution, with means equal to the published estimates of the stock, standard deviations

equal to our estimated standard errors of the stock, and correlation zero. We do not account for sampling uncertainty in the flow of new firearms q_t^N, as these quantities are measured without sampling uncertainty in administrative data. We form a standard error for $\hat{\varphi}$ as the standard deviation of our estimates across bootstrap replicates.

We perform inference on our estimates of the relative lethality of handguns χ_h via the standard sandwich GMM estimator of the asymptotic variance, treating each state-year as an observation that on expectation matches the handgun share of homicides as 0.899. We do not adjust for uncertainty in $\hat{\varphi}$, nor in the published estimates of the aggregate elasticity of homicide with respect to the firearm stock.

E.4 Robustness

We consider a number of alternative models of the homicide process and their implications for our key qualitative and quantitative conclusions. Briefly, each alternative model induces different quantitative relationships between firearms and homicides, as summarized in Table E.1. Figure E.1 displays the quantitative differences between these specifications in their predictions, though they all carry the same qualitative implications and policy recommendations as our preferred model in the main text. The remainder of this section provides detail.

E.4.1 Zero social cost from justifiable homicides

We consider the possibility that some share of the firearm homicides created by the firearm stock $Q_{c,s,t}$ do not generate social cost. In particular, the FBI categorizes justifiable homicides by private citizens as "The killing of a felon, during the commission of a felony, by a private citizen," which are "willful killings that must be reported as justifiable or excusable."[94] We remove these civilian justifiable homicides from our model by scaling down our predictions of expected firearm homicides $d_{c,s,t}$ by the approximate share of all U.S. firearm homicides that were not civilian justifiable firearm homicides over our sample period (3%), i.e. $(1 - 0.03)E[d_{c,s,t}]$.

As only a small share (3 percent) of civilian firearm homicides are considered "justifiable" by the FBI, this alternate assumption has only a small quantitative impact on our results. Our alternate assumption does not use information on justifiable homicides by peace officers ("the killing of a felon by a peace officer in the line of duty"), as our data on firearm fatalities from the NCHS explicitly exclude such incidents already.

[94] See: https://ucr.fbi.gov/crime-in-the-u.s/2019/crime-in-the-u.s.-2019/topic-pages/expanded-homicide

E.4.2 Distinct homicide elasticities

We reverse the assumption on the homicide process of Section 6, that the handgun and long gun stocks generate homicides through functions with the same shape $\kappa_h = \kappa_l = \kappa$, but with different scales $\chi_h \neq \chi_l$. Instead, we assume that these functions have different shapes $\kappa_h \neq \kappa_l$ but the same scales $\chi_h = \chi_l$. To calibrate these two elasticities, we use identical calibration targets and a process analogous to Section 6.2.

As handguns account for a much larger share of firearm fatalities than long guns, this alternate assumption produces $\kappa_h \gg \kappa_l$. The estimated handgun elasticity is about twice the size of the long gun elasticity. Relative to our baseline model, this assumption compresses differences in fatalities across firearm classes when the stock $Q_{c,s,t}$ is low and exacerbates differences in fatalities when the stock is high.

E.4.3 No convenience yield from handguns: equivalent assault rates

We consider alternate assumptions on the scale of the homicide functions: imposing that handguns and long guns generate equivalent risk of assault $\chi_h^a = \chi_l^a = \chi^a$, and that assaults from these two weapon classes differentially translate into fatalities based on differences in the inherent "technology" of these weapons $\chi_h^d \neq \chi_l^d$. Libby and Corzine (2007) provide an estimate of $\chi_h^d/\chi_l^d \approx 1.27$ by using FBI data to analyze conditional differences in the occurrence of fatalities among assaults perpetrated by firearms of different classes. Calibrating the ratio χ_h^d/χ_l^d to the Libby and Corzine (2007) estimates, we set κ to match the elasticity target from the prior literature (0.294), and calibrate $\zeta_{s,t}$ to match the average homicide levels in each state-year.

Heuristically, we refer to this assumption as removing the "convenience yield" from handguns or as specifying a "technology-only" homicide process. Our base model in Section 6 admits two potential reasons for the marginal handgun and long gun to generate different percentage changes in fatalities: either the marginal handgun generates greater assault risk ($\chi_h^a > \chi_l^a$, due to ease of carrying or concealment) and/or the marginal handgun assault is more lethal than the marginal long gun ($\chi_h^d > \chi_l^d$, due to ease of aim or differences in use cases). The estimates of Libby and Corzine (2007) imply that $\chi_h^a/\chi_l^a \gg \chi_h^d/\chi_l^d$, so that our alternate assumption of equal assault risk between handguns and long guns removes most of the differences in overall fatality risk between the classes χ_c.

This alternate assumption considerably compresses differences in homicides between the firearm classes at all levels of the firearm stock $Q_{c,s,t}$.

E.4.4 Used firearm transactions ω_c enter the homicide process

We consider an alternate model of the homicide process in which the cumulative purchase of new *and used* firearms affects fatalities. To implement this assumption, we alter the law of motion of the firearm stock to

$$\tilde{Q}_{c,s,t} = \begin{cases} (1-\varphi)\tilde{Q}_{c,s,t-1} + q^N_{c,s,t} + q_{\omega_c,s,t} & \text{if } t > 2015 \\ Q_{c,s,2015} & \text{if } t = 2015 \end{cases}$$

where $q_{\omega_c,s,t}$ is the quantity of used firearms of class c transacted in state s during year t.

This alternate law of motion captures the idea that, although firearms transactions among private individuals do not alter the overall level of firearm stock, they may alter the composition of firearm ownership, which may itself contribute to homicides (Rosenberg 2025). By holding the homicide process constant under this alternate law of motion, this alternate specification corresponds to the assumption that (i) the homicide risk from marginal purchasers of used firearms is exchangeable with marginal purchasers of new firearms, and that (ii) the sale of a used firearm does not alter the homicide risks generated by the seller.

As this assumption alters the effective stock of firearms $Q_{c,s,t} \neq \tilde{Q}_{c,s,t}$, we recalibrate the model's parameters using the procedure of Section 6.2, replacing $Q_{c,s,t}$ by $\tilde{Q}_{c,s,t}$. We do not adjust our calibration of the scrap rate φ, as this corresponds to the distribution of physical firearms in the economy, not their allocation across consumers and transactions.

This alternate assumption depresses the harm from a marginal firearm transaction, since $\tilde{Q}_{c,s,t} > Q_{c,s,t}$, and the expected count of firearm homicides experiences diminishing marginal returns. This alternate assumption attenuates the estimated public health impacts of our counterfactual taxes by around 40%, though the qualitative conclusions remain.

E.4.5 Harm from gun transactions depends on preferences, not chosen firearm class

Overview Our most substantive alternative model ties the harms from firearm transactions to the immutable preference characteristics of the consumers who transact them, not to the class of weapon actually chosen by these consumers under alternative policy regimes. Implementing this alternative model requires considerably more machinery than the others in this section. The benefit from this additional machinery is that we are able to apportion the harms from a firearm purchase to an individual, rather than to the weapon's class. Heuristically, the alternative model in this section corresponds to the idea that individuals who prefer handguns generate more harm from firearm purchases than individuals who prefer long guns, but there is no role for the cross-product differences in lethality to generate

that harm (i.e., "dangerous people prefer handguns, and they would be just as dangerous had they purchased a long gun").

To implement this idea, we track the choices of each consumer i under a policy environment τ, with $\tau = 0$ denoting the status quo of 10-11% excise taxes. Contrary to the model in Section 6, if a consumer would purchase a new handgun under $\tau = 0$ and a new long gun under $\tau = a$, their purchase counts towards the counterfactual flow of new handguns and not new long guns. We make the reverse assumption for consumers switching from new long guns to new handguns. We treat all other modes of substitution—to/from the outside option and used firearms—as in Section 6. Aside from adjusting these flows in the law of motion of the firearm stock, we leave the other parameters of the model $(\varphi, Q_{c,s,2015}, \kappa_c, \chi_h/\chi_l, \chi_l\zeta_{st})$ at the values used in Section 6.

Although simple to explain, this alternative model is challenging to implement for two separate reasons. A change in the policy τ may have impacts that are non-marginal and multi-dimensional (e.g., a finite tax increase on even a single product in a multi-product equilibrium). This means that common summaries of cross-product demands (e.g., diversion ratios) cannot properly summarize the full pattern of consumer substitution among the counterfactual policies we consider. Moreover, if i switches products from j to j' as the policy environment changes from $\tau = 0$ to $\tau = a$, this provides both the typical information used in the analysis of discrete choice models ($j \succeq_i j' \quad \forall j'$ under $\tau = 0$) and an additional piece of information ($j' \succeq_i j'' \quad \forall j''$ under $\tau = a$). Together, these preference relations provide information about the possible realizations of $\epsilon_{i,j,s,t}$ (the random utility component that varies across products within an individual and market), inducing posteriors on these unobservables that do not follow a standard distribution.

To address these challenges, we specialize the result of de Palma and Kilani (2011) to our specification of utility in Section 5. This allows us to derive exact, integral expressions for the relevant "switching" probabilities: of consumers from new handguns to new long guns (and vice versa), when the policy τ changes, and the market adjusts in equilibrium. Although lacking closed form, these probabilities are easy to approximate numerically. Their computation requires time similar to two evaluations of our constrained MLE objective of the demand model in Section 5, when all parameters including the demand shocks ξ_{jt} are fixed.

The remainder of this subsection provides details on the implementation of this robustness check. Equation (E.1) is the specification we implement.

Choice set and utilities Let $\mathcal{H}_t$ and $\mathcal{L}_t$ denote the sets of new handguns and long guns in period t, respectively, with $\mathcal{J}_t^N = \mathcal{H}_t \cup \mathcal{L}_t$. The set of used firearms is $\{\omega_h\} \cup \{\omega_l\}$. The

set of all firearms in the handgun class is $J_{h,t} = \mathcal{H}_t \cup \{\omega_h\}$. Allowing a no-purchase outside option, the choice set is $\mathcal{J}_t = \mathcal{H}_t \cup \omega_h \cup \mathcal{L}_t \cup \omega_l \cup \{0\}$.

We write the utility to consumer i under policy τ as

$$u_j^\tau = v_j^\tau + \zeta_0 + \rho_0 \zeta_{c(j)} + \rho_1 \tilde{\epsilon}_j,$$

where the final three terms are the two-level nested logit shocks described in Section 5. The first term v_j^τ captures the deterministic components of utility within a consumer-state-year-product (e.g., due to their individual-specific taste for particular product attributes). Throughout, we assume that the policy τ affects utilities only through this deterministic component v_j^τ and not through the nested logit shocks $(\zeta_0, \zeta_c, \tilde{\epsilon}_j)$.

Switching probabilities Integrating over the distribution of non-deterministic utility components, for an arbitrary vector of deterministic utility components $\vec{v}$, denote the probability that i chooses j by

$$P_{i,j}(\vec{v}) = P_i\left(\{v_j + \epsilon_{i,j} \geq v_{j'} + \epsilon_{i,j'}, \quad \forall j'\}\right)$$

where $\epsilon_{i,j}$ denotes the nested logit idiosyncratic shock. As the policy τ changes, we define the change in utility of alternative j by

$$\Delta_j \equiv v_j^{\tau=a} - v_j^{\tau=0},$$

with the utility change of the class-c used good Δ_{ω_c} and the outside option $\Delta_0 = 0$. The vector of changes in deterministic utility across all options is $\vec{\Delta}$.

We also denote the probability that consumer i switches from choice j under the status quo $\tau = 0$ to k under the alternative $\tau = a$ by

$$\begin{aligned} P_{i,j\to k}^a &= P_i\left(\left\{u_{ij}^{\tau=0} \geq u_{ij'}^{\tau=0};\ \text{all } j'\right\} \cap \left\{u_{ik}^{\tau=a} \geq u_{ik'}^{\tau=a};\ \text{all } k'\right\}\right) \\ &= P_i\left(\left\{u_{ij}^{\tau=0} \geq u_{ij'}^{\tau=0};\ \text{all } j'\right\} \cap \left\{u_{ik}^{\tau=0} + \Delta_{i,k} \geq u_{ik'}^{\tau=0} + \Delta_{i,k'};\ \text{all } k'\right\}\right), \end{aligned}$$

such that $\sum_j \sum_k P_{i,j\to k}^a = 1$.

Then, the result of de Palma and Kilani (2011) implies that as the policy τ changes from

0 to a, holding the vector of nested logit shocks $\vec{\epsilon}_{i,s,t}$ fixed:

$$P^a_{i,j\to k} = \begin{cases} 0 & \text{if } \Delta_k < \Delta_j \\ P_{i,j}\left(\vec{v}_i^{\tau=0} + \max\left\{\vec{\Delta} - \Delta_j, \vec{0}\right\}\right) & \text{if } j = k \\ \int_{\Delta_j}^{\Delta_k} -\frac{\partial P_{i,j}\left(\vec{v}_i^{\tau=0} + \max\left\{\vec{\Delta} - z, 0\right\}\right)}{\partial v_{i,k}^{\tau=0}} dz & \text{if } \Delta_k > \Delta_j, \end{cases}$$

where the maximum is taken component-wise over the vector of alternatives. The final line is the total change in the probability that i would choose alternative j, caused by the utility change in alternative k spanned by the shift in policy from $\tau = 0$ to $\tau = a$.

At any arbitrary vector of non-random utilities $\vec{v}$, it is a standard application of logit arithmetic to show that the partial derivative is

$$-\frac{\partial P_{i,j}(\vec{v})}{\partial v_k} = P_{i,j}(\vec{v}) \times P_{i,k}(\vec{v}) \times \begin{cases} 1 & \text{if } c(j) \neq 0 \text{ and } c(k) = 0 \text{ (or reversed)} \\ \left(1 + \frac{1/\rho_0 - 1}{P_{i,\text{in}}(\vec{v})}\right) & \text{if } c(j) \neq c(k) \text{ and } c(j), c(k) \neq 0 \\ \left(1 + \frac{1/\rho_0 - 1}{P_{i,\text{in}}(\vec{v})} + \frac{1/\rho_1 - 1/\rho_0}{P_{i,\text{in}}(\vec{v}) \cdot P_{i,c|\text{in}}(\vec{v})}\right) & \text{if } c(j) = c(k) \text{ and } c(j), c(k) \neq 0. \end{cases}$$

where the choice probabilities are, respectively, the probability i purchases a new or used firearm (i.e., not the outside option) and the probability that i purchases a new or used firearm of class c, conditional on not choosing the outside option:

$$\begin{aligned} P_{i,\text{in}}(\vec{v}) &= 1 - P_{i,0}(\vec{v}) = 1 - \left(1 + \exp\left\{\rho_0 I_{i,0}\right\}\right)^{-1} \\ P_{i,c|\text{in}}(\vec{v}) &= \exp\left\{\frac{\rho_1}{\rho_0} I_{i,c} - I_{i,0}\right\}, \end{aligned}$$

and these terms depend on the inclusive values:

$$\begin{aligned} I_{i,c} &= \log \sum_{j:c(j)=c} \exp\left\{\frac{v_j}{\rho_1}\right\} \\ I_{i,0} &= \log \sum_{c'} \exp\left\{\frac{\rho_1}{\rho_0} I_{i,c'}\right\}. \end{aligned}$$

This characterizes the probability that i switches from j to k as the policy changes from the status quo $\tau = 0$ to an alternative $\tau = a$.

Quantities under the alternate model Using the above, we can express the mass of consumers each state-year that switch their purchase from a new handgun under the status

quo to a new long gun under the alternative (the object of interest):

$$q^a_{s,t,h\to l} = \sum_d M_{s,d} \int_{\vec{\nu}} \left(\sum_{j\in\mathcal{H}_t} \sum_{k\in\mathcal{L}_t} P^a_{i,j\to k}(D_i = d, t, \nu_i = \nu, s) \right) \varphi(\vec{\nu}) d\vec{\nu}.$$

To evaluate this quantity, we use quadrature twice: once to approximate the integral in the switching probability inside the double sum, conditional on the random coefficients $\vec{\nu}$; and again to evaluate the mixture of these approximated switching probabilities across the distribution of random coefficients $\varphi(\cdot)$. Similar expressions hold for switching probabilities across all collections of products.

Finally, the alternate handgun and long gun purchase quantities we feed into the law of motion of the firearm stock are, respectively,

$$\check{q}^a_{h,s,t} = q^a_{h,s,t} + q^a_{s,t,h\to l} \tag{E.1}$$

$$\check{q}^a_{l,s,t} = q^a_{l,s,t} - q^a_{s,t,h\to l} \tag{E.2}$$

We then simulate the predictions of our health model assuming Equation E.1 governs the flow of new firearms.

Section	Name	Description	Change from Base Model
E.4.1	Remove CJH	No welfare cost from civilian justifiable homicides, tracked by FBI	$VSL = 0.97 \times \$9.5M$
E.4.2	Distinct elasticities	Different curvature linking gun stock to homicides for handguns and long guns. Functions constrained to same scale.	$\kappa_h \neq \kappa_l$ and $\chi_h = \chi_l$
E.4.3	Equal assault rates	Handguns are no more likely to be used in assaults than long guns. Homicide scales by class from "instrumentality" $P_c(\text{die}\|\text{shot})$	$\chi_h/\chi_l = 1.27$ (Libby and Corzine 2007)
E.4.4	Used quantity	Firearm flow in each class's law of motion includes used goods ω_c	$q_{c,s,t} = q_{\omega_c,s,t} + \sum_{j\in\mathcal{J}:c(j)=c} q_{j,s,t}$
E.4.5	Switchers quantity	Counterfactual firearm flow in the handgun law of motion includes consumers switching from new handguns to new long guns	$q_{h,s,t} = \sum_{j\in\mathcal{J}:c(j)=h} q_{j,s,t} + \sum_{j,k\in\mathcal{J}:c(j)=h,c(k)=l} q_{j\to k,s,t}$

Table E.1: Summary of alternative health models

Table displays the roles of alternative assumptions on homicide model of Section 6, described in Section E.

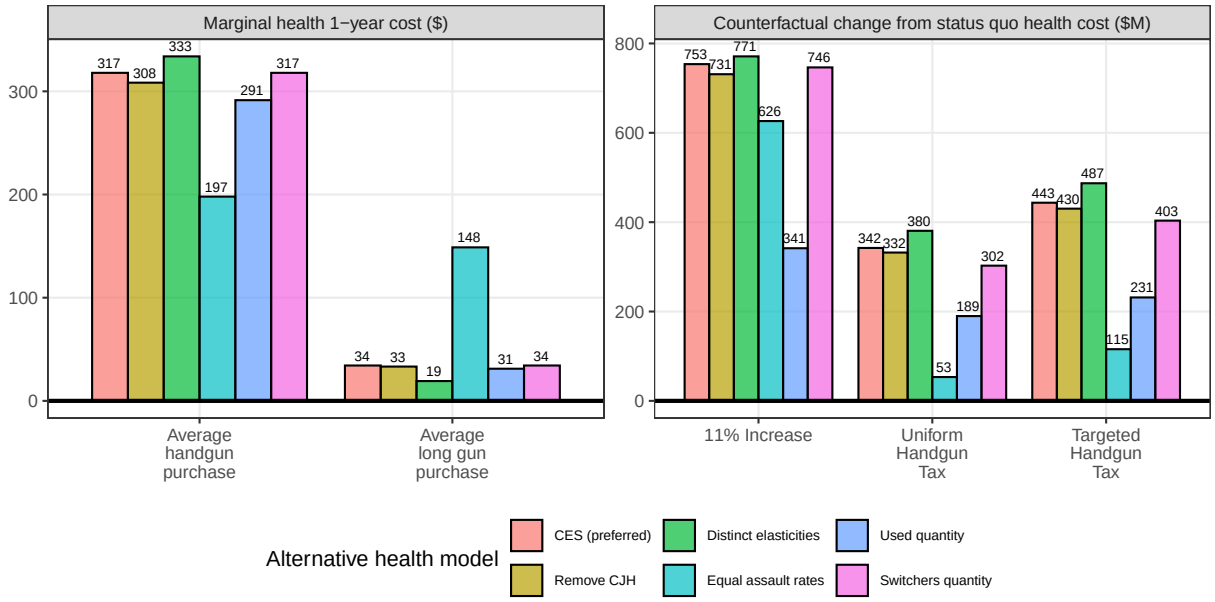


Figure E.1: Changes in marginal and counterfactual homicide costs under different health models

Figure displays the changes to our health model's predictions under the alternative assumptions summarized in Table E.1 and described in Section E. The left-hand plot displays different values of the average marginal homicide cost of a new gun purchase during our sample, analogous to values reported in Table 4. The right-hand plot displays different values of the change in homicide cost under different counterfactual taxes, analogous to values reported in Figure 3. Colors represent alternative assumptions on the homicide process, with red (CES Preferred) denoting our baseline model.

F Details on Counterfactuals

This section provides details on our analysis of counterfactual firearm regulations in Section 7.

F.1 Welfare components

This section formalizes the components appearing in our measure of welfare as a function of a counterfactual tax change $\Delta\vec{v}$.

Implementing the counterfactual tax change $\Delta\vec{v}$ leads to a counterfactual tax rate $(\vec{v} + \Delta\vec{v})$. Firms re-optimize under the counterfactual tax rate, setting yearly prices $\vec{p}_t(\Delta\vec{v})$. Consumers re-optimize their purchases—facing the new prices $\vec{p}(\Delta\vec{v})$—giving rise to counterfactual yearly quantities $\vec{q}_t(\Delta\vec{v})$.

Industry profits in year t are

$$\Pi_t(\Delta\vec{v}) = \sum_f \Pi_{f,t}(\Delta\vec{v}) = \sum_{j\in\mathcal{J}_{f,t}} (1 - v_j - \Delta v_j) \times p_{j,t}(\Delta\vec{v}) \times q_{j,t}(\Delta\vec{v}) - c_{j,t} \times q_{j,t}(\Delta\vec{v}), \quad \text{(F.1)}$$

where the summation is across profits from each firm f during year t. The in-sample change in profits is

$$\Delta\Pi(\Delta\vec{v}) = \sum_{t=2016}^{2022} \Pi_t(\Delta\vec{v}) - \Pi_t(\vec{0}). \quad \text{(F.2)}$$

Consumer surplus in state s during year t is calculated in its typical form, as expected utility divided by the price coefficient:

$$CS_{s,t}(\Delta\vec{v}) = M_s \int_{i:s(i)=s} \frac{1}{\alpha_i} E_i \left[\max_{j\in\mathcal{J}_t} u_{i,j,s,t}(\vec{p}_t(\Delta\vec{v})) \right] \partial F_s(i), \quad \text{(F.3)}$$

where the utilities depend on the tax change $(\Delta\vec{v})$ indirectly through the prices $\vec{p}(\Delta\vec{v})$. The integral is taken across the joint distribution of observed and unobservably heterogeneous consumers in state s. The inner expectation can be computed under the assumed nested logit distribution on the i.i.d. shocks $\epsilon_{i,j,s,t}$. The in-sample change in consumer surplus is

$$\Delta CS(\Delta\vec{v}) = \sum_{t=2016}^{2022} \sum_s CS_{s,t}(\Delta\vec{v}) - CS_{s,t}(\vec{0}). \quad \text{(F.4)}$$

We also use the value

$$CS_0 = \sum_{t=2016}^{2022} \sum_s CS_{s,t}(\vec{0}) \quad \text{(F.5)}$$

to calculate the consumer surplus constraint for the constrained optimal taxes.

Tax revenue to the federal government from the sales of firearms from manufacturers in year t is

$$G_t(\Delta\vec{v}) = \sum_j (v_j + \Delta v_j) \times p_{j,t}(\Delta\vec{v}) \times q_{j,t}(\Delta\vec{v}) \tag{F.6}$$

The in-sample change in government revenue is

$$\Delta G(\Delta\vec{v}) = \sum_{t=2016}^{2022} G_t(\Delta\vec{v}) - G_t(\vec{0}). \tag{F.7}$$

We do not consider other forms of revenue, such as state and local sales taxes, corporate taxes, or income taxes on the workers who manufacture firearms.

The in-sample change in homicides in state s and year t is

$$\Delta d_{s,t}(\Delta\vec{v}) = \sum_c d_{c,s,t} \left[\left(1 + \frac{\sum_{t'=2016}^{t} (1-\varphi)^{t-t'} \left(q_{c,s,t'}(\Delta\vec{v}) - q_{c,s,t'}(\vec{0}) \right)}{(1-\varphi)^{t-2015} Q_{c,s,2015} + \sum_{t'=2016}^{t} (1-\varphi)^{t-t'} q_{c,s,t'}(\vec{0})} \right)^{\kappa} - 1 \right], \tag{F.8}$$

which increases over time t. The overall in-sample change in homicides is

$$\Delta d(\Delta\vec{v}) = \sum_s \sum_{t=2016}^{2022} \Delta d_{s,t}(\Delta\vec{v}). \tag{F.9}$$

F.2 Optimal Tax Implementation

We do not attempt to estimate or implement the full solution to the planner's constrained-optimal tax problem in Propositions 1 or 2. From Equation (A.2), computing these rates would require estimates of the marginal externality ϕ_j, the profit margin μ_j, the monopoly margin μ_j^M, the shadow price of the consumer surplus constraint λ, and the price p_j. Each of these is an equilibrium object, which will vary across the fundamentals of different markets t and the actual tax policy in effect $\vec{v}$. Computing the planner's solution would therefore require a complex search across a high-dimensional space of tax systems $\vec{\tau} \in \mathbb{R}^{\sum_t ||\mathcal{J}_t^N||}$ for the location of a fixed point.

Instead, we implement an estimable approximation to the planner's solution, denoted the "Targeted Handgun Tax" in our counterfactuals. We compute this approximation under a sufficient statistics approach (e.g., Allcott et al. 2019), by evaluating the right-hand side of Equation (A.2) using estimates of the relevant objects under a natural, counterfactual tax system that may not be a fixed point. We take the approach below to estimating each of the terms in the targeted handgun tax.

- $\hat{\tau}_l = 0$: To ease computation, we set all counterfactual taxes on long guns to 0. Figure 2 shows that long guns have profit margins and negative externalities of similar magnitude—approximately canceling in the optimal tax formula—while the negative externalities from handguns dominate all other welfare components.

- $\hat{\mu}_j$: To estimate the equilibrium profit margin, we plug in the equilibrium profit margins that obtain under our counterfactual uniform handgun tax.

- $\hat{\mu}_j^M$: To estimate the monopoly profit margin, we plug in the margins that obtain under the counterfactual monopoly ownership of all new firearms at the baseline excise taxes of 10%-11%, as in Figure F.3.

- $\hat{\lambda}$: We use the fact that, in equilibrium, the shadow cost of our political feasibility constraint is $\lambda = -\partial W/\partial CS_0$, with CS_0 the minimal consumer surplus that satisfies the constraint. To estimate this shadow cost, we calculate the in-sample changes to consumer surplus and social welfare that would obtain if the counterfactual uniform handgun tax were raised by 0.1 percentage points (as would happen if the political feasibility constraint were slightly more slack). Then, we calculate $\hat{\lambda}$ as the ratio of the change in aggregate welfare to the change in consumer surplus from this small tax increase:
$$\hat{\lambda} = -\frac{W(v_h = 15.6, v_l = 0) - W(v_h = 15.5, v_l = 0)}{CS(v_h = 15.6, v_l = 0) - CS(v_h = 15.5, v_l = 0)} = 1.8$$
Thus, relative to the optimal uniform handgun tax that allows no drop in counterfactual consumer surplus, this estimate implies that allowing consumer surplus to drop by 1 dollar would enable a re-optimized uniform handgun tax to generate 1.8 additional dollars of social welfare.

- $\hat{\phi}_h$: We ease computation by assuming that the marginal externality is identical for each new handgun in each state during each year of our sample $\phi_{j,s,t} = \phi_{j',s',t'} = \phi_h$. To compute this value, we use the fact that the average handgun purchase during our sample generated an annual marginal harm from firearm homicides of $MSC_h = 0.00003347 \times \$9.5M \approx \$318$, each year it remained in circulation. We also use the fact that, during our 7-year sample, the average handgun purchase would be observed in the handgun stock over approximately 4 years (i.e., our sample covers 2016–2022, so the average handgun is purchased in roughly 2019, and would be observed in the stock from 2019–2022). Combining these facts, we estimate the marginal externality as
$$\hat{\phi}_h = MSC_h \cdot \frac{1 - (1 - \varphi)^4}{1 - (1 - \varphi)}$$

where $\varphi = 0.015$ is the scrap rate of firearms from the stock. This approximation also matches the structure of the in-sample welfare objective by not discounting firearm homicides that occur later in the sample.

With these estimates, we start by estimating an approximation to the optimal specific tax as:

$$\hat{\tau}_{j,t} = \hat{\phi}_h - \hat{\mu}_{j,t} - \hat{\lambda} \cdot \hat{\mu}^M_{j,t} \tag{F.10}$$

We then make a few alterations to ensure the tax scheme is comparable across years and implementable empirically.

We impose that the average targeted handgun tax does not drift over time, easing interpretation over our medium-run sample. For instance, removing drift ensures that the tax system does not drastically shift up or down each year, due to over-time changes in market fundamentals (e.g., a temporary demand shock or a change in market power from a merger). In reality, these changes in fundamentals may be hard to forecast, and to condition a tax system on, so we view our assumption of no drift as an approximation to the policy design problem facing regulators. Removing drift also precludes changes in the tax system due to discounting (e.g., setting high taxes in early years to prevent homicides in the short term).

To remove drift, we demean the initial calculation of the tax in Equation (F.10) by year, netting out differences attributable to year-specific changes. Call this version of the tax $\tilde{\tau}_{j,t}$. Finally, to ensure the tax can be implemented as an ad valorem tax, we scale the tax to be between 0 and 1 and set the final tax as:

$$v_{j,t} = g \cdot \frac{\tilde{\tau}_{j,t} - \min_{j,t}\{\tilde{\tau}_{j,t}\}}{\max_{j,t}\{\tilde{\tau}_{j,t}\} - \min_{j,t}\{\tilde{\tau}_{j,t}\}} \tag{F.11}$$

where g is a hyperparameter chosen to ensure that the resulting tax system meets our definition of political feasibility.

To compute the targeted handgun tax under the assumption of perfect competition, we follow the exact same procedure, but modify Equation (F.10) by zeroing out the equilibrium margin μ_j (i.e., $\hat{\tau}_{j,t} = \hat{\phi}_h - \hat{\lambda} \cdot \hat{\mu}^M_{j,t}$). Removing the margin reflects the design of the optimal specific tax under perfectly competitive supply in Proposition 1.

F.3 Benchmarking Against Estimated Effects of Existing Gun Policies

In this section, we detail our approach to benchmarking the public health effects of our counterfactual taxes, against the effects of existing policy interventions in the public health literature.

Our basic strategy is to take all studies that examine the effect of each type of firearms regulation, filter to studies that are most credible in terms of their empirical design, and aggregate across studies to produce an average estimated treatment effect. We start with a database of studies from RAND's *Science of Gun Policy* project (Smart et al. 2024), which compiles hundreds of studies that leverage state-level variation in firearm policies to estimate the effect of these policies on public health outcomes. The database standardizes each published estimate of a state-level gun policy's effect into a logged incidence rate ratio (IRR): the ratio of the outcome's rate in an adopting state to the rate that state would have experienced absent the policy. This can be interpreted as the percent change in an outcome.

From these studies, we retain estimates that match four criteria: (i) RAND's quality review judges free of serious methodological concerns (final designations of "minor" concerns or "high quality"); (ii) use firearm homicides in the general population as the outcome, excluding subcategories such as intimate partner homicides or homicides of juveniles; (iii) report an IRR with a 95% confidence interval, from which we recover the standard error of the log IRR as the width of the log-scale confidence interval divided by 2×1.96; and (iv) analyze policies that have at least 2 studies that match the prior three criteria. The underlying studies are predominantly state-year panel regressions (negative binomial, Poisson, or log-linear) of homicide rates on a policy indicator with state and time fixed effects, so most IRRs measure the average proportional change in homicides across a study's post-adoption state-years.

From the set of studies of each firearm law, we estimate a pooled treatment effect (IRR) using the weighted average of study-level estimates, with weights proportional to the inverse variance of each estimate (Cochran 1954). The implied standard error for the pooled IRR estimate is $1/\sqrt{\sum_k 1/\hat{\sigma}_k^2}$, where k indexes studies, and σ_k the associated standard error on the IRR estimate from that study.

For our own estimates, we convert the values from Section 7 to their corresponding IRR. This is calculated as the predicted number of firearms homicides under each counterfactual tax policy, relative to the status quo equilibrium. We calculate these at the state-level, and aggregate by taking a weighted average across states, with weights proportional to observed firearms homicides. This weighting choice parallels the inverse-variance weighting used to compute our pooled RAND estimates: as homicides in our model follow a Poisson-distribution, homicide weights correspond exactly to precision weights.[95] We report this pooled IRR in 2022, the final year of our sample, allowing each policy to accumulate in strength through 7 years of implementation. This horizon is the closest in-sample analog to the horizon of the RAND estimates, where the average post-treatment period is typically greater than 10 years. If anything, our own estimates are conservative relative to the horizon

[95] In practice, unweighted and population-weighted averages produce very similar estimates.

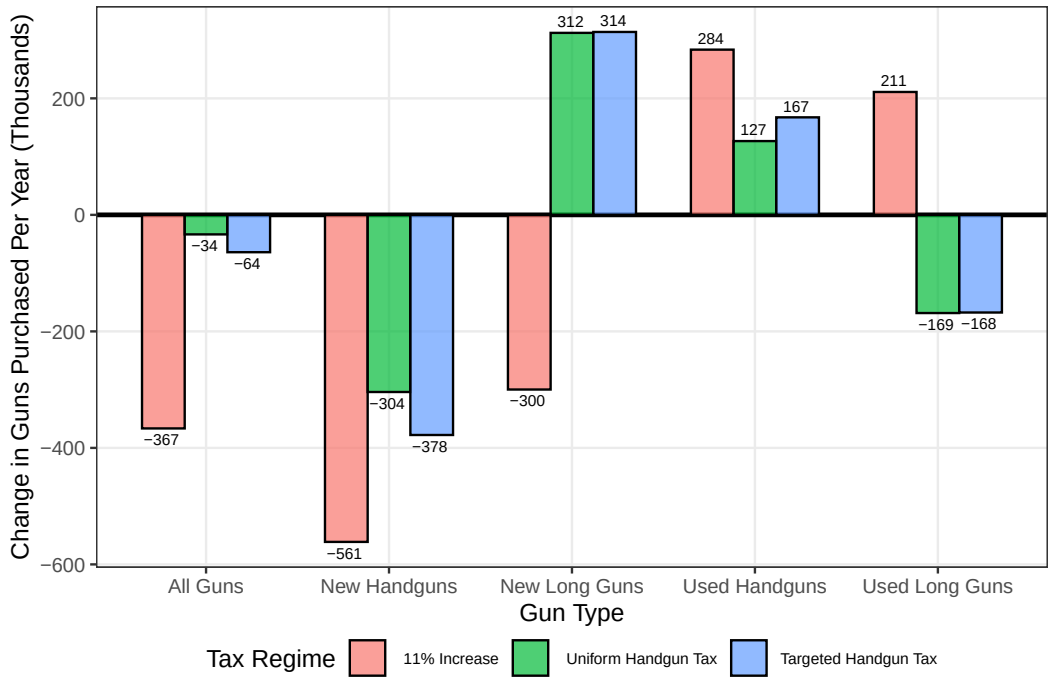
of the comparison studies.

Table F.1 reports the estimates from the meta-analysis against our own estimates. The more restrictive gun laws from the RAND meta-analysis (Panel A) are estimated to reduce firearm homicides by 1.7%–10.1%. As a comparison, our uniform handgun taxes that are designed to hold consumer surplus neutral only reduce firearm homicides by 0.3 to 0.4%. Of the policies covered in Smart et al. (2024), the pooled estimates for implementing Background Checks or Waiting Periods, and removing Right-to-Carry or Stand-Your-Ground laws are statistically significant at conventional levels. Our estimates suggest that reformulating the tax schedule to target handguns would prevent 3%-6% as many homicides as these laws.

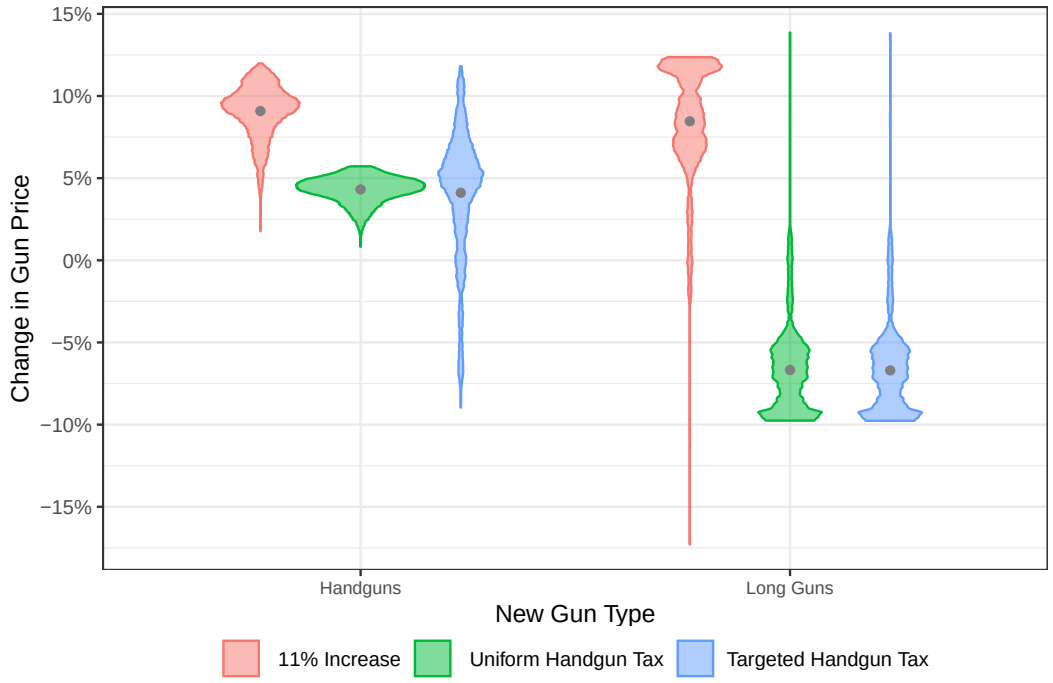
Policy	Log(IRR)	SE	# Studies
Panel A: Existing Policies (RAND Meta-Analysis)			
Background Checks	-0.1011	0.0278	6
Child Access Prevention Carry	-0.0172	0.0342	2
Minimum Age Requirements	-0.0148	0.0152	6
Right-to-Carry	0.0566	0.0149	8
Stand-Your-Ground	0.0563	0.0141	6
Waiting Periods	-0.0917	0.0277	3
Panel B: Counterfactual Firearms Taxes (2022)			
Double Tax	-0.0070	—	—
Uniform Handgun Tax	-0.0031	—	—
Targeted Handgun Tax	-0.0040	—	—

Table F.1: Comparison of Firearm Tax Policies to Existing Interventions

Table displays changes in firearm homicide fatalities caused by the gun policy interventions, collected in Smart et al. (2024), relative to our estimated counterfactual tax policies. Changes are reported as the policies effects on the log Incidence Rate Ratio, such that 0 represents no effect. Panel A reports fixed-effects meta-analytic estimates and standard errors for the change in the log(IRR), pooling the quasi-experimental literature on existing state gun policies, across the #Studies results compiled by Smart et al. (2024). Panel B reports the homicide-weighted mean across states of the state-level change in the log(IRR) implied by each counterfactual tax policy in 2022, relative to the observed equilibrium under status quo taxes. Appendix F.3 provides further details on the calculations in this table.



(a) Change in Purchases by Gun Type



(b) Change in Prices

Figure F.1: Effects of Firearms Tax Policies on Prices and Quantities

Figure shows changes in prices and quantities caused by the counterfactual firearm taxes we consider. In Panel (a), we show the aggregate change in gun sales in the average year, for different types of firearms. In Panel (b), we display a violin plot of the distribution of the change in prices among firearm model-years under active production, separately by firearm class The dot in each density violin plot represents the unweighted average percentage change in prices.

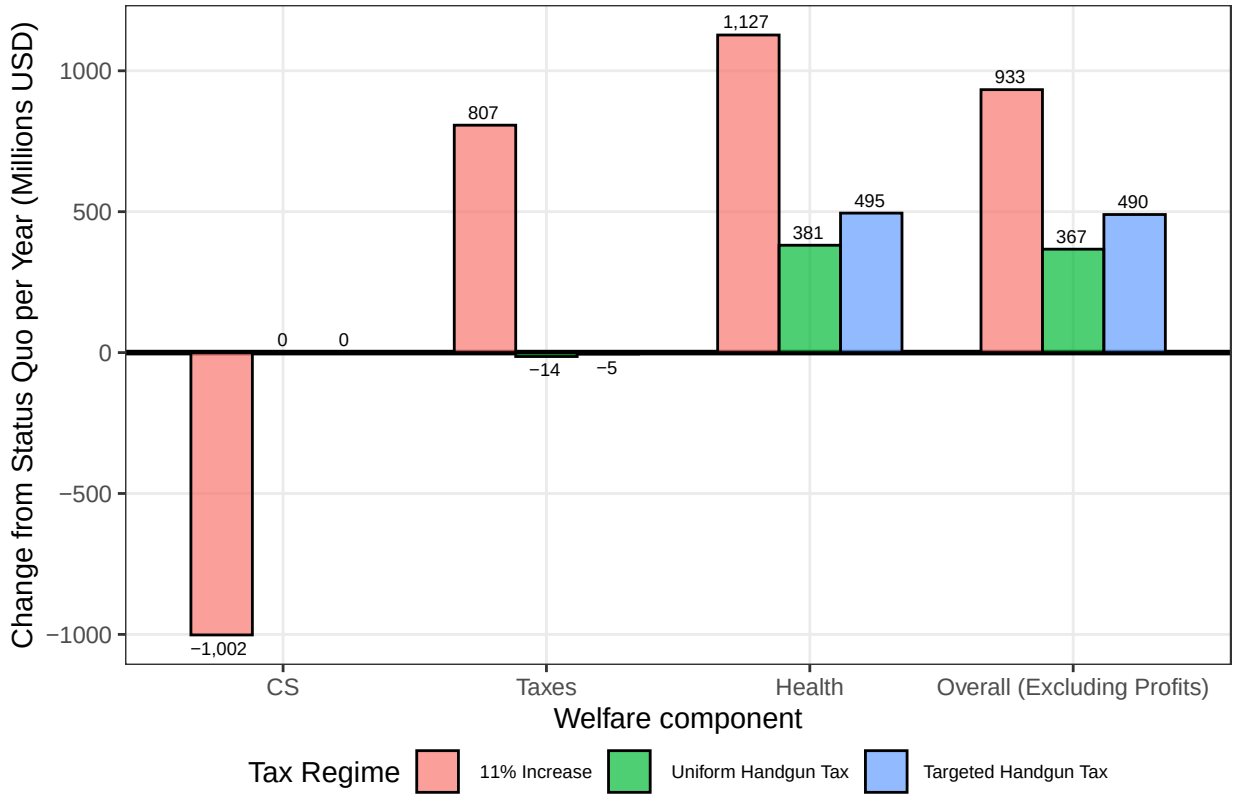


Figure F.2: Welfare Effects of Firearms Tax Policies Under Competitive Pricing

Figure shows the change in dollars of welfare nationwide, comparing the status quo to the counterfactual firearm taxes in Section 7, under the alternative assumption that manufacturers make no profits in equilibrium $p_{j,t} = c_{j,t}/(1 - v_{c(j)})$. Counterfactuals tax systems: (orange) homogeneously raise the tax on all firearms by 11 percentage points, (green) set the tax on all long guns to 0% and on all handguns to 15.5%, (blue) set the tax on all long guns to 0% and a set heterogeneous tax on each handgun-year by approximating the constrained-optimal tax in Equation (3). Vertical axis shows the average in-sample annual change in each of the four welfare components from Equation (1) and the change in overall welfare (their sum).

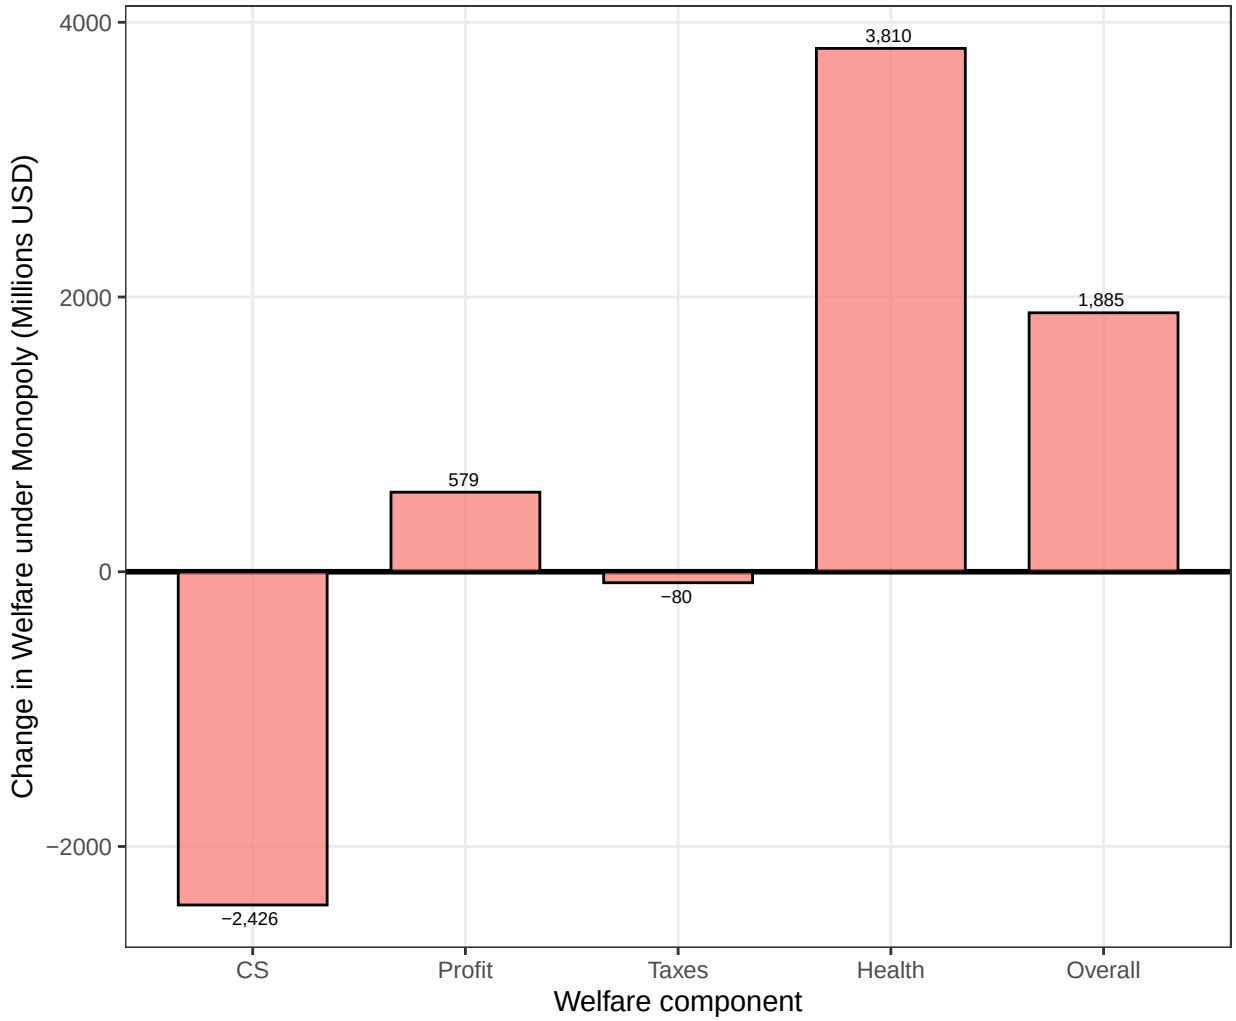


Figure F.3: Welfare Effects of Monopoly Pricing

Figure shows the change in dollars of welfare nationwide, comparing the status quo to prices that would be set under counterfactual static-monopoly ownership of all firearm models in active production each year. Vertical axis shows the average in-sample annual change in each of the four welfare components from Equation (1) and the change in overall welfare (their sum). Monopoly pricing conduct is characterized in Appendix A.

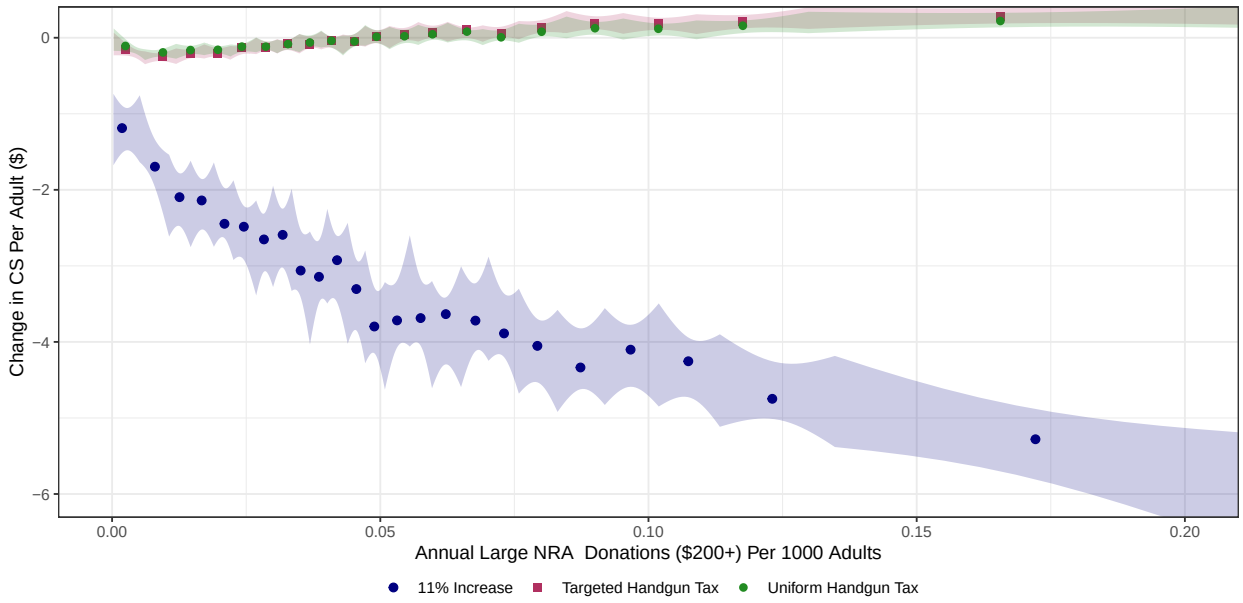


Figure F.4: Effect of Tax Reforms on Consumer Surplus, by NRA Prominence (Donations Per Capita)

Figure is a binned scatterplot across U.S. congressional districts, showing the change in a measure of average consumer surplus relative to NRA donations, for each of the counterfactual tax systems we consider. Vertical axis is the counterfactual change in the average annual dollar value of consumer surplus (CS) per consumer in each congressional district, averaged over 2016-2022. Horizontal axis is the average annual number of donations $\geq$ \$200 to the NRA, per 1,000 consumers between 2016-2022. Counterfactuals tax systems: (blue) homogeneously raise the tax on all firearms by 11 percentage points, (green) set the tax on all long guns to 0% and on all handguns to 15.5%, (red) set the tax on all long guns to 0% and a set heterogeneous tax on each handgun-year by approximating the constrained-optimal tax in Equation (3). Consumer surplus is calculated analogously to Equation (F.3). The procedure of Cattaneo et al. (2019) is used to select the number and location of points and to form 95% confidence intervals.