Firearm Policy Buyers Would Accept*

Bradley T. Shapiro[†] Sara Drango[‡] Sarah Moshary[§]

April 2026

Abstract

Policies restricting gun sales face substantial political opposition and constitutional constraints. We take a different approach: rather than asking whether gun restrictions *should* be imposed, we ask how much it would *cost* to compensate potential handgun buyers to choose alternative products. Using a structural model of firearms demand estimated from stated-choice experiments, we estimate the distribution of willingness-to-accept (WTA) for substituting away from handguns toward long guns or no gun purchase. Because handguns account for the majority of firearm deaths in the United States, substitution away from handguns may have meaningful public health implications. We find that targeted compensation can induce substantial substitution at relatively modest cost, while uniform subsidies are far less cost-effective due to windfall payments to inframarginal buyers. Our results highlight scope for compensation-based policies to reduce externalities while leaving buyers better off.

JEL Codes: D12, I18, K42, H23

Keywords: Firearms, gun control, willingness to accept, demand estimation, public policy

*We thank seminar participants at The University of Florida, Johns Hopkins, Stanford and The University of Texas for helpful comments. ChatGPT and Gemini were used to edit and refine prose. All errors are our own.

[†]University of Chicago Booth School of Business. Email: bradley.shapiro@chicagobooth.edu

[‡]University of Chicago Booth School of Business.

[§]University of California, Berkeley Haas School of Business.

1 Introduction

Approximately 45,000 people die from firearm injuries each year in the United States, including roughly 20,000 homicides and 24,000 suicides.[1] Handguns account for a disproportionate share of homicides, representing approximately 90% of homicides where firearm type is identified.[2] Evidence suggests that handgun stocks in communities have persistent effects on violence (Evans et al., 2022). At the same time, policies that restrict gun ownership face intense political opposition. Proposed regulations routinely fail in Congress, and even modest restrictions mobilize well-organized advocacy (Goss, 2006; Lacombe, 2021). This opposition reflects not only organized interests, but also strong preferences: estimated willingness-to-pay for firearms is high (Moshary et al., 2025; Rosenberg, 2025). As a result, policies that restrict firearm access impose substantial political and consumer costs.

We examine the problem from a different perspective. Rather than asking whether gun restrictions *should* be imposed, we consider policies that compensate consumers for substituting toward less harmful alternatives. In the context of firearms, this means estimating how much it would cost to compensate potential handgun buyers to choose a long gun or forgo purchase.

Switching away from handguns may reduce firearm violence; handguns are easier to conceal, easier to use in impulsive acts of violence, and given their short barrel length and light weight, more conducive to self-harm. Compensation-based substitution policies are feasible here because viable alternatives exist: most firearm owners state personal and household safety as their primary motivation for ownership (Stantcheva et al., 2025), and long guns (rifles and shotguns) can serve this protective function.

Using a structural model of firearm demand, we estimate the distribution of willingness-to-accept (WTA) of handgun buyers to substitute away from handguns toward long guns or no purchase. We focus on gun-naïve households—those without existing firearms—because

[1] CDC WONDER database, 2022 data.

[2] FBI Uniform Crime Reports, 2019–2022.

the literature identifies causal effects of gun ownership at the extensive margin (Duggan, 2001; Cook and Ludwig, 2006). In contrast, the existing stock of roughly 400 million firearms is largely fixed in the short run. We use elasticity estimates from the literature to simulate how model-predicted changes in gun ownership might translate into reduced mortality. Because this relationship is uncertain, we consider alternative assumptions: an optimistic scenario in which only handguns affect homicide rates and a pessimistic scenario in which long guns contribute proportionally to mortality. We then compute the total cost of alternative subsidies and compare the implied cost per life saved to standard estimates of the value of a statistical life.

Our approach is motivated by an asymmetry central to the political economy of gun regulation. The costs of gun ownership are diffuse: potential victims cannot predict who will bear them, which makes collective action difficult (Olson, 1965). In contrast, the benefits are concentrated and salient: gun owners know what they would lose from restrictions. Compensation-based policies change these incentives by making gun owners weakly *better off*. A person who accepts a subsidy to choose a long gun over a handgun reveals that the subsidy exceeds their switching cost. Compensation-based approaches have proven politically viable in other domains, including agricultural subsidies, cash-for-clunkers programs, and tax credits for electric vehicles (Goulder and Parry, 2008).

Constitutional considerations provide another reason to consider compensation-based policies. Following *District of Columbia v. Heller* (2008) and subsequent decisions, courts have interpreted the Second Amendment to impose meaningful constraints on firearm regulation (Cook et al., 2013; Blocher and Miller, 2018). Courts have struck down outright bans on common firearms categories. Compensation-based policies may face fewer constitutional obstacles because they encourage alternatives while leaving individuals free to purchase any legal firearm.

We consider two hypothetical policy scenarios that benchmark the cost of substituting away from handguns. In Scenario A (perfect targeting), the social planner observes each prospective handgun buyer's willingness-to-accept and offers exactly the minimum payment to induce substitution. This represents the minimum cost of achieving any given reduction in handgun purchases. In Scenario B (uniform subsidy), the planner cannot observe individual preferences and instead offers the same subsidy to everyone who purchases a long gun or chooses the outside option. This scenario is more realistic but more costly because it generates substantial windfall payments to inframarginal buyers.

Our analysis is not a policy blueprint but a benchmarking exercise. By comparing the distribution of willingness to accept for substituting away from handguns to the associated mortality benefits, we assess whether compensation-based approaches could be cost-effective in this setting. The goal is not to specify an implementation, but to establish whether this class of policies merits further consideration. If so, questions of targeting, mechanism design, and administration become important next steps.

Our estimates reveal that the distribution of compensation required to induce substitution away from handguns is highly skewed. The median prospective handgun buyer requires approximately \$812 to switch, while the mean exceeds \$2,000. Modest payments can therefore induce many buyers to substitute, but reaching the most committed handgun buyers is expensive.

Our findings imply that compensation-based policies could be highly cost effective. Under perfect targeting, the optimal subsidy ranges from \$435 to \$959, depending on whether increases in long gun ownership also contribute to homicides. These subsidies avert 68–275 deaths per year at a total program cost of \$0.5–\$1.5 billion. The implied \$5–\$7 million cost per averted death is well below the \$13 million value of statistical life commonly used in regulatory analysis. A targeted subsidy therefore passes cost-benefit analysis even under pessimistic assumptions about the relationship between long gun ownership and violence.

Under uniform subsidies, estimated cost-effectiveness depends critically on elasticity assumptions. Under optimistic assumptions, a \$229 uniform subsidy averts 100 deaths at \$12.5 million per life. Under pessimistic assumptions, however, *no positive uniform subsidy passes* *cost-benefit analysis*. The minimum achievable cost per life is approximately \$17 million, well above the VSL threshold. This difference arises because uniform subsidies generate large windfall payments to inframarginal buyers; when the marginal response is small, these windfalls dominate the policy's benefits.

Information about consumer preferences is therefore extremely valuable. At a given number of switchers, uniform subsidies cost 5–6 times as much as targeted subsidies. At the optimal uniform subsidy, only 26% of program expenditure goes to consumers who actually switch; the remaining 74% are windfall payments to inframarginal recipients. This pattern underscores the potential value of screening mechanisms or partial targeting.

Our paper makes several contributions. First, we provide the first quantitative analysis of compensation-based firearm policies. While political scientists have documented the obstacles to gun control in the United States (Goss, 2006; Lacombe, 2021) and recent industrial organization research evaluates taxation (Armona and Rosenberg, 2025) and retail access (Rosenberg, 2025), we study whether compensating prospective handgun buyers to choose less lethal alternatives could reduce morality at reasonable cost. In contrast to policies based on restrictions or taxation, compensation-based approaches explicitly account for the political and constitutional constraints that shape firearm regulation.

Second, we bridge the guns-and-violence literature (Duggan, 2001; Cook and Ludwig, 2006) with demand estimation. Prior work establishes that gun ownership increases homicides but says little about the determinants of firearm demand and how ownership might be altered via policy to reduce mortality. By estimating the distribution of willingness-to-accept for substituting away from handguns, we connect these literatures and quantify the costs of inducing buyers to switch.

Finally, we extend the demand framework of Moshary et al. (2025) to counterfactual policy analysis. Leveraging the model's treatment of consideration set formation, we compute willingness-to-accept for substitution across firearm categories and evaluate the cost-effectiveness of compensation-based policies. More broadly, the willingness-to-accept distribution we estimate provides an empirical input for evaluating any incentive-based approach to reducing handgun purchases. Our framework treats consumers' stated motivations for firearm ownership, particularly concerns about personal safety (Stantcheva et al., 2025), as legitimate preferences and asks what compensation would be required to shift behavior.

The remainder of the paper proceeds as follows. Section 2 describes the stated-choice experiment used to elicit firearm preferences. Section 3 presents the search cost demand model and discusses identification. Section 4 reports the estimated model's predictions for market shares in the status quo. Section 5 develops the counterfactual analysis, including the translation from individual to household market shares, the mortality model, and results for both policy scenarios. Section 6 discusses implications and concludes.

2 Data: Stated Choice Experiment

Our analysis uses data from a stated-choice experiment designed to elicit preferences over firearms. This section describes the experimental design, sample characteristics, and key features of the data. For complete details, see Moshary et al. (2025).

2.1 Experimental Design

Sample. The experiment was fielded online in March and April 2022 to 22,522 U.S. adults in partnership with The Harris Poll. Survey participants were drawn from the general population available to The Harris Poll. Among respondents, 4,018 report owning a gun, interest in buying a gun in the next 12 months, and/or interest in purchasing a gun more generally. This set comprises the estimation sample.

Choice Tasks. Each respondent completed seven choice tasks. In each task, respondents were presented with three firearm alternatives and asked, in a forced-choice hypothetical purchase scenario, to select their most preferred option. Respondents were then asked whether they would actually buy their most preferred alternative at the quoted price. This design requires respondents to consider relative trade-offs among options rather than making absolute judgments, and the follow-up purchase question allows us to observe the extensive margin.

Attributes. In each choice task, respondents were shown pictures of each alternative firearm, alongside information on price, brand, caliber, capacity, and action type. We estimate part-worths for category (Pistol, Revolver, Rifle, Shotgun), brand (e.g., Smith & Wesson, Glock, Remington), and price.

Randomization. For each task, prices were varied independently of other attributes, allowing identification of price sensitivity. The composition of choice sets—the three gun models available in each task—was also randomized, providing variation to help identify consideration set formation.

2.2 Consideration Sets

A key feature of the experimental design is that it elicits both choices and *consideration sets*. Before the choice tasks, respondents report which types of firearms they would consider among pistols, revolvers, shotguns, and/or rifles. The alternatives in each task were then drawn primarily from these stated categories but tasks occasionally included options outside the reported consideration set. This design allows us to observe directly whether a respondent considers handguns, long guns, or both. It also lets us identify preferences over firearms outside their consideration set, which is necessary for computing willingness-to-accept for switching to non-considered categories. These data also provide direct evidence on how buyers search before purchasing.

2.3 Sample Characteristics

Table 1 reports summary statistics for the respondents. The sample is roughly balanced by gender and somewhat older than the general population, consistent with the demographics of gun ownership in the United States. Approximately two-thirds of respondents own firearms.

Table 1: Sample Characteristics

	Mean
Male	0.48
Age	48.8
Household income <\$50k	0.42
Household income \$50k–\$100k	0.37
Household income >\$100k	0.21
College educated	0.31
Currently owns a gun	0.64
No gun in household	0.18

The distinction between respondents from gun-owning households and those from gun-naïve households is central to our analysis. Policies that shift marginal purchases toward long guns leave the extensive margin of handgun ownership unchanged among households that already own a handgun, whereas the literature linking gun ownership to homicide largely focuses on that margin. The mortality implications of changing the *type* of a marginal purchase within a gun-owning household are therefore less clear than those of changing *whether* a gun-naïve household acquires a handgun.

We define gun-naïve households as those in which the respondent reports no current gun ownership. Approximately 8.45% of U.S. households are gun-naïve prospective buyers: they are considering a firearm purchase but do not currently own a firearm.[3]

[3] In our sample, approximately 19.2% of participants reside in households with no firearm. Given that our sample frame is current owners or those considering buying a firearm and the survey data shows a relatively constant ≈44% of households own at least one firearm year over year, we infer that $0.44 \cdot 0.192 \approx 0.0845$ – roughly 8.45% of total U.S. households are gun-naïve and considering a purchase.

2.4 Design Features and Limitations

Stated choice experiments offer several advantages for studying firearm demand. They generate rich variation in prices and attributes, allowing us to observe substitution patterns that would be difficult to recover from transaction data. The two-stage design—selecting a preferred alternative, then deciding whether to purchase—allows us to identify the extensive margin. Eliciting consideration sets directly provides information on the search process. Survey data are also well-suited here because firearm purchases are sensitive and administrative data are limited.

The main limitation of stated choice data is that choices are hypothetical. Respondents may not fully account for budgets, storage, or other real-world constraints. We address this concern in two ways. First, the experimental design presents realistic scenarios with plausible prices. Second, Moshary et al. (2025) validated stated preferences to revealed preference data where available. More broadly, the literature on hypothetical bias suggests that it primarily affects absolute willingness-to-pay rather than marginal rates of substitution between attributes (Carlsson and Martinsson, 2001). Because our counterfactual analysis focuses on relative trade-offs—how much additional compensation is required to induce substitution from handguns to long guns—the estimates are likely more robust to hypothetical bias than absolute valuations.

2.5 Summary

The stated choice experiment provides rich data on firearm preferences among prospective buyers. The key features for our analysis are: (1) choices across handguns, long guns, and an outside option; (2) exogenous price variation for identifying demand elasticities; (3) consideration set information for identifying search costs; and (4) the ability to distinguish gun-naïve households from existing gun owners. The next section describes how we use these data to estimate a structural model of firearm demand.

3 Model: Demand Estimation with Consideration Data

This section presents the structural model of firearm demand developed in Moshary et al. (2025). Our data include respondents' consideration sets in addition to their choices. To incorporate this information, we adopt a search cost framework in which consideration sets are determined endogenously. This structure allows the model to match both observed choices and consideration patterns, improving identification of preference parameters.

3.1 Model Setup

Products and Categories. The market consists of J products organized into K categories. In our application, the four categories are pistol, revolver, rifle, and shotgun. Products are specific firearm profiles defined by brand, price, and feature combinations. The outside option is always available and requires no search cost. Let $\mathcal{K} = \{1, \ldots, K\}$ denote the set of categories and $\mathcal{J}_k$ denote the set of products in category k.

Consumer Problem. Consumer i makes two sequential decisions:

1. **Consideration**: Choose which categories to consider, incurring search cost κ_i per category.
2. **Choice**: Choose the utility-maximizing option from their consideration set (including the outside option).

Timing. The timing is as follows:

1. Consumer with prior beliefs about category-level utilities observes their search cost κ_i.
2. Consumer chooses consideration set $\mathcal{C}_i \subseteq \mathcal{K}$, paying κ_i for each category they include.
3. Consumer observes the utilities of all products in considered categories.

4. Consumer chooses the utility-maximizing option from $\mathcal{C}_i \cup \{0\}$, where 0 denotes the outside option.

3.2 Utility Specification

Product Utility. Consumer i's utility from purchasing product j in category k is:

$$u_{ij} = \alpha_{ik} + x_j'\beta_i + \gamma_i p_j + \varepsilon_{ij} \tag{1}$$

where α_{ik} is consumer i's preference for category k, x_j is a vector of product attributes (brand dummies, features), β_i is consumer i's vector of attribute preferences, p_j is the price of product j, $\gamma_i < 0$ is consumer i's price sensitivity, and ε_{ij} is an idiosyncratic taste shock distributed Type I Extreme Value. We allow for heterogeneity in preferences across consumers. Individual preference parameters $(\alpha_i, \beta_i, \gamma_i, \kappa_i)$ are drawn from a joint distribution in the population:

$$(\alpha_i, \beta_i, \gamma_i, \kappa_i) \sim F(\cdot|\theta) \tag{2}$$

where θ parametrizes the distribution (means and the variance–covariance matrix).

Outside Option. The outside option has utility:

$$u_{i0} = \varepsilon_{i0} \tag{3}$$

where ε_{i0} is also distributed Type I Extreme Value. The outside option is normalized to have zero mean utility.

Category-Level Expected Utility. Before searching a category, consumer i does not observe the product-level shocks $\{\varepsilon_{ij}\}_{j \in \mathcal{J}_k}$. The expected purchase utility from category k, conditional on consideration, is:

$$V_{ik} = \mathbb{E}\left[\max_{j \in \mathcal{J}_k} u_{ij}\right] \tag{4}$$

Under the Type I Extreme Value assumption, this has a closed-form expression:

$$V_{ik} = \log\left(\sum_{j\in\mathcal{J}_k} \exp\left(\alpha_{ik} + x_j'\beta_i + \gamma_i p_j\right)\right). \tag{5}$$

This is the "inclusive value" of category k for consumer i.

3.3 Consideration Set Formation

Optimal Consideration. Consumer i chooses consideration set $\mathcal{C}_i$ to maximize expected utility net of search costs:

$$\max_{\mathcal{C}\subseteq\mathcal{K}} \left\{ \mathbb{E}\left[\max_{j\in\bigcup_{k\in\mathcal{C}}\mathcal{J}_k\cup\{0\}} u_{ij}\right] - |\mathcal{C}|\cdot\kappa_i \right\}. \tag{6}$$

The first term is the expected utility from the best option in the consideration set. The second term is the total search cost, which is proportional to the number of categories considered.

Solving the Consideration Problem. Under the Type I Extreme Value assumption, the expected maximum utility from consideration set $\mathcal{C}$ is:

$$\mathbb{E}\left[\max_{j\in\bigcup_{k\in\mathcal{C}}\mathcal{J}_k\cup\{0\}} u_{ij}\right] = \log\left(1 + \sum_{k\in\mathcal{C}} \exp(V_{ik})\right). \tag{7}$$

The optimal consideration set is defined as:

$$\mathcal{C}_i^* = \arg\max_{\mathcal{C}\subseteq\mathcal{K}} \left\{ \log\left(1 + \sum_{k\in\mathcal{C}} \exp(V_{ik})\right) - |\mathcal{C}|\cdot\kappa_i \right\}. \tag{8}$$

This is a combinatorial optimization problem with 2^K possible consideration sets. For our application with $K = 5$ categories, enumeration is feasible.

Intuition. A category is included in the consideration set if its marginal inclusive value exceeds the associated search cost. Categories with high preference (α_{ik}) or many attractive products are more likely to be considered. Consumers with low search costs (κ_i) consider more categories. Although the model includes search costs, they are not the primary objects of interest. Rather, incorporating search allows the model to match observed consideration patterns, facilitating inference about the underlying preference parameters.

3.4 Choice Probabilities

Conditional on Consideration Set. Given consideration set $\mathcal{C}_i$, the probability that consumer i chooses product j in category $k \in \mathcal{C}_i$ is:

$$P(j|\mathcal{C}_i) = \frac{\exp(u_{ij})}{\exp(u_{i0}) + \sum_{k' \in \mathcal{C}_i} \sum_{j' \in \mathcal{J}_{k'}} \exp(u_{ij'})}. \tag{9}$$

This is the standard multinomial logit formula, conditional on the consideration set.

Unconditional Choice Probability. The unconditional probability of choosing product j integrates over the distribution of consideration sets:

$$P(j) = \sum_{\mathcal{C}: k(j) \in \mathcal{C}} P(\mathcal{C}) \cdot P(j|\mathcal{C}) \tag{10}$$

where $k(j)$ is the category containing product j and $P(\mathcal{C})$ is the probability that consideration set $\mathcal{C}$ is optimal, which is a function of search costs and preference parameters.

3.5 Identification

The model parameters are identified by different sources of variation in the data.

Price Sensitivity (γ_i). Price sensitivity is identified by the randomized variation in prices across choice tasks. Holding other attributes fixed, we observe how choice probabilities change as prices vary. The coefficient γ_i measures the marginal utility of money and is essential for translating utility differences into dollar-denominated willingness-to-accept.

Attribute Preferences (β_i). Preferences over product attributes (brand, features) are identified by randomized variation in the attributes present in each choice set. For example, comparing choices across tasks that differ solely in the brand of one product reveals the preference for that brand.

Category Preferences (α_{ik}). Category preferences are identified by choices among products within and across categories, holding other attributes fixed. The relative frequency of choosing handguns vs. long guns (conditional on both being in the consideration set) identifies the category preference parameters.

Search Cost (κ_i). The search cost is identified by variation in consideration sets, *holding fixed the inclusive values of categories.* Identification comes from comparing consumers who make similar choices when presented with the same alternatives but differ in their stated consideration sets. For example, consider two consumers with the same inclusive values $\{V_{ik}\}$ but different consideration sets $\mathcal{C}_i \neq \mathcal{C}_{i'}$. Conditional on preferences, differences in consideration behavior identify differences in search costs; consumers who consider more categories (given the same preferences) have lower search costs.

3.6 Estimation

We estimate the model using hierarchical Bayesian methods implemented in the `bayesm` package (Rossi et al., 2005). The estimation uses Markov Chain Monte Carlo (MCMC) with a Metropolis-in-Gibbs sampler, following the approach in Moshary et al. (2025).

The hierarchical structure treats individual-level preference parameters $(\alpha_i, \beta_i, \gamma_i, \kappa_i)$ as draws from a population distribution characterized by hyperparameters θ. The Gibbs sampler alternates between: (1) drawing individual-level parameters conditional on the hyperparameters and data, using a Metropolis-Hastings step to account for the nonlinear consideration set model; and (2) drawing hyperparameters conditional on the individual-level parameters. We use the posterior draws to simulate the distribution of willingness-to-accept for each consumer, which we use in the counterfactual analysis. For additional estimation details, see Appendix A.

3.7 Estimation Results Summary

We estimate the model using the full sample of 4,018 respondents. The estimation yields posterior distributions for individual-level parameters $(\alpha_i, \beta_i, \gamma_i, \kappa_i)$ as well as population-level hyperparameters. Table 2 reports summary statistics for the key parameters.

Table 2: Estimated Model Parameters

Parameter	Mean	Median	5th Pctl	95th Pctl
Price (per $100)	−0.34	−0.16	−1.25	−0.02
Handgun intercept	0.56	0.59	−3.55	4.66
Long gun intercept	0.05	0.09	−3.64	3.72
Search cost ($)	420	126	9	1,659

Notes: Posterior distribution summaries across all respondent-draw combinations. Category intercepts are simple averages across subcategories (4 for handguns, 10 for long guns). Search cost is converted to dollars using each draw's price coefficient. See Moshary et al. (2025) for additional details.

Table 3 reports the same statistics for the gun-naïve subsample, which is the focus of our counterfactual analysis.

Three results are particularly relevant for our counterfactual analysis. First, consumers are price-sensitive: the median price coefficient of −0.16 per $100 implies that a $100 price increase reduces utility by 0.16 units. Second, there is substantial heterogeneity in preferences for handguns vs. long guns, with many consumers strongly preferring one type over another. Third, search costs are meaningful, with a median of roughly $126. This implies that inducing substitution toward long guns requires sizable subsidies for a subset of consumers who do

Table 3: Estimated Model Parameters: Gun-Naïve Subsample

Parameter	Mean	Median	5th Pctl	95th Pctl
Price (per $100)	−0.38	−0.19	−1.41	−0.02
Handgun intercept	0.61	0.59	−3.40	4.65
Long gun intercept	−0.05	−0.09	−3.67	3.69
Search cost ($)	395	110	8	1,575

Notes: Posterior distribution summaries for gun-naive respondents only. Gun-naive households exhibit slightly higher price sensitivity and a stronger preference for handguns relative to long guns.

not consider rifles or shotguns in the status quo; the subsidy must be large enough to shift both product choice and, when relevant, the consideration margin.

4 Model Results: Status Quo Market Shares

This section reports predicted firearm market shares in the status quo. We present results for both the full sample and the gun-naïve subsample, which is the focus of our counterfactual analysis.

4.1 Simulation Procedure

To generate market share predictions, we simulate choices from the estimated model. For each respondent i, we:

1. Draw preference parameters from the posterior distribution: $(\alpha_i, \beta_i, \gamma_i, \kappa_i)$.

2. Compute the optimal consideration set $\mathcal{C}_i^*$ given these parameters.

3. Draw logit shocks ε_{ij} for all products in $\mathcal{C}_i^*$ and the outside option.

4. Determine the utility-maximizing choice.

We repeat this process for multiple posterior draws to integrate over parameter uncertainty. The resulting choice distribution provides market share predictions that account for both preference heterogeneity and estimation uncertainty.

4.2 Market Shares

Table 4 reports predicted market shares for the full sample of prospective gun buyers. Handguns account for over half of predicted purchases, with pistols being the dominant subcategory. Long guns account for about one-third of purchases. Approximately 7% of prospective buyers choose the outside option—that is, they consider purchasing but ultimately do not buy.

Table 4: Predicted Market Shares: Full Sample vs. Gun-Naïve

Category	Full Sample	Gun-Naive
Handguns	**58.2%**	**63.6%**
Pistol	38.3%	43.0%
Revolver	19.9%	20.6%
Long Guns	**34.8%**	**28.7%**
Semi-auto rifle/shotgun	12.2%	10.0%
Other rifle	10.1%	7.4%
Other shotgun	12.6%	11.4%
Outside Option	**7.0%**	**7.7%**

We highlight two differences in predicted market shares between gun-naïve and gun-owning households. First, gun-naïve households have a higher predicted handgun share compared to the full sample (63.6% vs. 58.2%). This may reflect that first-time buyers are more likely to prioritize home defense or personal protection, uses for which handguns are often preferred. Gun-naïve households also have a higher outside option share (7.7% vs. 7.0%), consistent with first-time buyers being more price-sensitive and less committed to purchasing than experienced gun owners..

4.3 Consideration Sets

Table 5 reports the distribution of predicted consideration sets among gun-naïve handgun buyers (i.e., simulated consumers whose utility-maximizing choice is a handgun). For all handgun buyers, inducing a switch to long guns requires compensation for selecting a less-preferred option. However, whether the buyer requires compensation for incurring incremental search costs varies. Among handgun buyers, 57.7% consider long guns even absent a long gun subsidy and so require no additional compensation. The remaining 42.3% consider only handguns. Among these consumers, a subsidy might entice some to consider long guns exclusively, changing which categories are searched, but not total search costs. Others might expand their search inducing them to switch requires compensating them not only for selecting a less-preferred option, but also potentially for expanding their search.

Table 5: Gun Categories in Consideration Set Among Gun-Naive Handgun Buyers

Gun Categories Considered	Share
Handguns only	42.3%
Handguns + Long guns	57.7%

Note: The outside option (no purchase) is always available and requires no search cost.

This distinction has important implications for the distribution of willingness-to-accept. Consumers who already consider long guns have lower WTA (the utility difference divided by price sensitivity), while many of those who do not consider long guns have higher WTA because the subsidy must also induce search.

4.4 Translation to Household-Level Ownership

Status Quo Household Ownership. To apply causal estimates of firearm ownership on mortality, we convert our individual-level market share estimates (Table 4) into predicted household firearm ownership. We construct status quo household ownership shares by gun type as follows. Our conjoint estimates imply 8.45% of households are gun-naïve; we treat the remaining gun-owning households as incumbent gun owners whose ownership predates the current period. Gallup reports that 44% of U.S. households contain at least one gun (Gallup, 2021), implying that incumbent households comprise roughly 35.55% of all households. We allocate incumbent households between handgun owners and long-gun-only owners using ownership ratios from the 2021 National Firearms Survey, yielding baseline shares of approximately 29.5% handgun households and 6.2% long-gun-only household.. We then add the gun-naive households' predicted choices from our demand model to obtain the status quo ownership shares reported in Table 6:

Table 6: Status Quo Household Gun Ownership

Household Type	Share of Households
Handgun household	38.7%
Long-gun-only household	4.8%
No gun	56.5%
Any gun	43.5%

In our policy counterfactuals, we connect simulated changes in individual purchases to changes in ownership shares as follows. Let s_{HG}^{gn} denote the share of gun-naïve households choosing handguns. A policy that reduces this share by Δs changes household-level handgun ownership by:

$$\Delta(\text{HG households}) = f_{gn} \cdot \Delta s \tag{11}$$

where $f_{gn} = 0.0845$ is the fraction of households that are gun-naïve. For example, reducing the handgun share among gun-naïve households from 63.6% to 50% implies a 1.15 pp reduction of the household-level handgun ownership share.

4.5 Summary and implications for counterfactuals

The estimated model implies that 63.6% of gun-naïve households would choose handguns in the status quo, quantifying the margin targeted by our counterfactual policies. The estimates also reveal substantial heterogeneity in consideration sets; some handgun buyers already consider long guns and can be induced to switch at relatively low cost, while others must be compensated to even consider the long gun category.

The next section uses the predicted market shares and substitution patterns to evaluate the cost and mortality effects of compensation-based policies.

5 Counterfactual Analysis

This section evaluates compensation-based policies that reduce handgun purchases. We first define willingness to accept for substitution away from handguns. We then present a mortality model that translates changes in gun ownership into averted deaths. We use these components to compare two policy scenarios: perfect targeting and uniform subsidies.

5.1 Willingness to Accept

Willingness to accept (WTA) is the minimum payment required to induce a household that would choose a handgun in the status quo to instead choose a less lethal alternative, either a long gun or the outside option. We adopt an ex post perspective, assuming that consumers have already formed their consideration sets prior to the introduction of the subsidy; they then decide whether to expand those sets (if applicable) and whether—and what—to purchase. In the appendix, we also estimate counterfactuals under an ex ante perspective, in which consumers learn about the subsidy before forming their consideration sets. We view the ex ante case as potentially more relevant for long-run policy. Buyers currently in the market (our sample) have likely already considered some alternatives, and the policy cannot reverse that exposure.

First, consider handgun buyers whose status quo consideration sets include long guns. Let U_i^{HG} denote consumer i's utility from their preferred handgun, U_i^{LG} their utility from their preferred long gun option, and U_i^{out} their utility from the outside option. A subsidy s

induces consumer i to switch away from handguns if either:

$$U_i^{LG} - \gamma_i \cdot s > U_i^{HG} \quad \text{(switch to long gun)} \tag{12}$$

$$U_i^{out} - \gamma_i \cdot s > U_i^{HG} \quad \text{(switch to outside)} \tag{13}$$

where $\gamma_i < 0$ is price sensitivity (so -$\gamma_i \cdot s > 0$ for subsidies). For these consumers, the minimum subsidy to induce a switch to a long gun is simply the monetized utility gap:

$$\text{WTA}_i^{LG} = \frac{U_i^{HG} - U_i^{LG}}{-\gamma_i} \quad \text{if LG} \in \mathcal{C}_i. \tag{14}$$

If long guns are not in the consumer's status quo consideration set, inducing substitution to a long gun requires inducing consideration of at least one long gun category. A subsidy s on non-handgun alternatives shifts the inclusive value of each long gun category k from V_{ik} to $V_{ik} + (-\gamma_i) \cdot s$. This has two effects: it raises the net utility of each long gun option, and in so doing, it can induce the consumer to evaluate one or both long gun categories. We compute this WTA numerically as the minimum subsidy such that (a) the consumer finds it worthwhile to add at least one long gun category into their consideration set at subsidized inclusive values and (b) the best long gun in the expanded consideration set, evaluated at subsidized utilities, exceeds the chosen handgun.

The minimum subsidy to induce a switch to the outside option is:

$$\text{WTA}_i^{out} = \frac{U_i^{HG} - U_i^{out}}{-\gamma_i}. \tag{15}$$

The outside option is always in the consideration set, so no additional search is required.

The consumer's overall WTA to switch away from a handgun purchase is the minimum of the two:

$$\text{WTA}_i = \min\left\{\text{WTA}_i^{LG}, \text{WTA}_i^{out}\right\}. \tag{16}$$

The consumer switches to whichever alternative requires the smaller subsidy. Let $d_i \in$

$\{\text{LG}, \text{out}\}$ denote the switch destination:

$$d_i = \begin{cases} \text{LG} & \text{if WTA}_i^{LG} < \text{WTA}_i^{out} \\ \text{out} & \text{otherwise.} \end{cases} \tag{17}$$

Our estimates imply that 78.9% of switchers would go to long guns and 21.1% would go to the outside option.

5.2 WTA Distribution

Figure 1 plots the distribution of WTA among gun-naïve handgun buyers and Table 7 reports summary statistics. The distribution is highly right-skewed: the median consumer requires \$812 to switch, while the mean requires \$2,111 due to a long right tail. This skewness implies that many consumers can be induced to switch away from handguns at modest cost, but reaching the most committed handgun buyers is expensive.

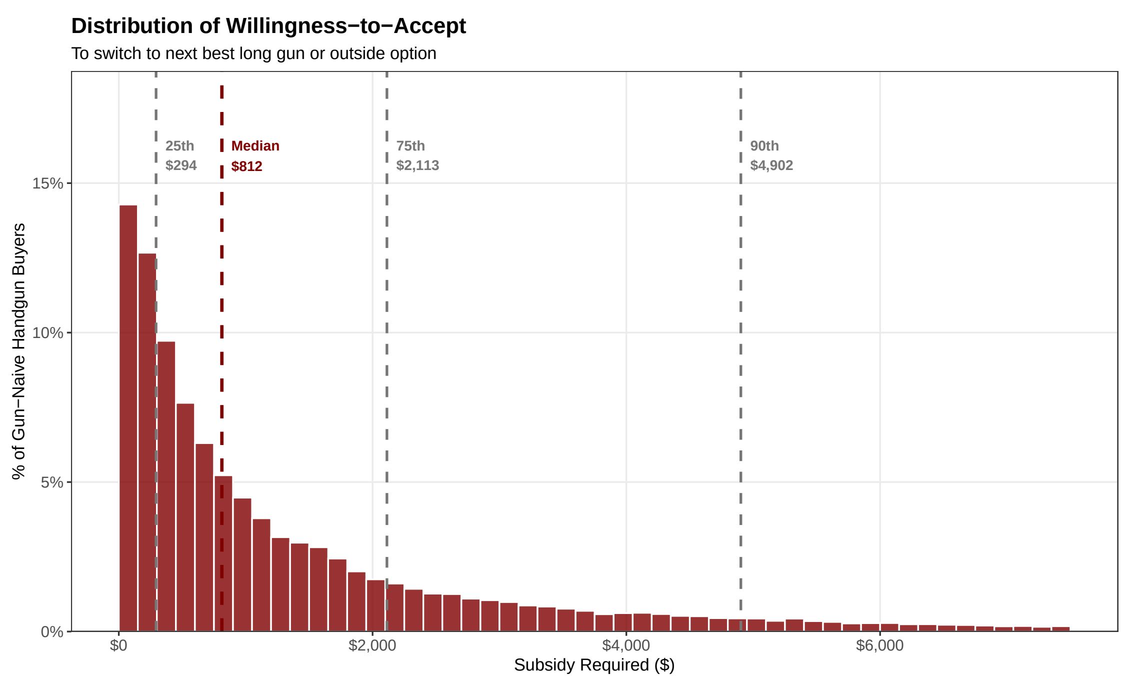

Figure 1: Distribution of Willingness to Accept Among Gun-Naïve Handgun Buyers

Table 7: WTA Summary Statistics

Statistic	Value
Mean	\$2,111
Median	\$812
5th percentile	\$57
25th percentile	\$294
75th percentile	\$2,113
95th percentile	\$8,040

5.3 From Ownership to Homicides: The Mortality Model

To translate changes in gun ownership into averted deaths, we model homicides as a function of gun ownership with constant elasticity:

$$H = H_0 \cdot \left(\frac{G}{G_0}\right)^{\varepsilon} \tag{18}$$

where H denotes homicides, G the share of households owning a firearm, subscript 0 denotes status quo values, and ε is the elasticity of homicides with respect to gun ownership. This specification implies that a 1% increase in gun ownership leads to an ε% increase in homicides.

We calibrate the elasticity using estimates from the economics and public health literatures. Prior work, including Duggan (2001), places the elasticity in the range of 0.1 to 0.3. These elasticity estimates are identified from extensive-margin variation in legal household gun ownership, the margin our policy targets. Duggan (2001) proxies for gun prevalence using magazine subscriptions, capturing changes in ownership driven by legal acquisition rather than criminal markets. Our counterfactual likewise focuses on gun-naïve households making first-time purchases, so these estimates are appropriate for our setting. We adopt $\varepsilon = 0.2$ as our central value, consistent with Cook and Ludwig (2006).

Our analysis focuses on homicides as a measure of harm. Gun ownership may generate additional externalities—including non-fatal injuries and psychological trauma—that are serious and consequential but not captured in our model. These outcomes are important in their own right and could be incorporated into a similar framework with appropriate data.

A natural concern is whether policies targeting legal first-time buyers can meaningfully affect homicide rates, given the association between gun violence and criminal activity. However, most homicides do not involve strangers: 76% of female and 56% of male victims are killed by intimate partners, family members, friends, or acquaintances, while only 12% of women and 10.6% of men were killed by strangers (Bureau of Justice Statistics, 2022). The relevant mechanism is therefore not substitution by criminals, but rather changes in the weapons available within households. The type of weapon on hand affects the lethality of violent incidents arising from domestic disputes, mental health crises, or interpersonal conflicts. This "instrumentality" effect, whereby weapon lethality shapes the probability of death, is well documented in the public health literature (Braga and Cook, 2018). Legal first-time buyers are precisely the population for whom household weapon choice matters.

Because handguns account for the majority of homicides, consistent with their ease of concealment and use in close-quarters confrontations, we allow for separate elasticities by firearm type:

$$H_{HG} = H_{HG,0} \cdot \left(\frac{G_{HG}}{G_{HG,0}} \right)^{\varepsilon_{HG}} \tag{19}$$

$$H_{LG} = H_{LG,0} \cdot \left(\frac{G_{LG}}{G_{LG,0}} \right)^{\varepsilon_{LG}} \tag{20}$$

where H_{HG} and H_{LG} denote handgun and long gun homicides, and G_{HG} and G_{LG} denote handgun and long-gun-only household shares.

Two Scenarios. We consider two scenarios that bracket a range of plausible effects, as summarized by Table 8. The optimistic scenario reflects the view that handguns are disproportionately associated with homicide risk, while long guns contribute relatively little. The pessimistic scenario instead assumes that substitution across weapon types does not meaningfully reduce homicide risk, and therefore assigns the same elasticity to long guns and handguns.

Table 8: Elasticity Scenarios

Scenario	ε_{HG}	ε_{LG}	Interpretation
Optimistic	0.22	0	Handgun reductions save lives; long gun increases have no effect
Pessimistic	0.2	0.2	Both gun types contribute proportionally to homicides

Our mortality model interprets the disparity between handgun and long gun homicide rates as reflecting, at least in part, a *treatment* effect of weapon type. This interpretation is consistent with the substantial physical differences between handguns and long guns types and with evidence that many homicides arise from escalated interpersonal conflicts, where weapon type can determine whether an incident becomes fatal (Cook and Ludwig, 2000).[4]

At the same time, our estimates rely on this maintained assumption. If differences in homicide rates reflect selection into weapon type rather than causal effects of the weapon itself, our approach would overstate the mortality benefits of substitution. In the extreme, if long guns were equally or more lethal for the relevant marginal buyers, substitution could yield little reduction in homicides or even increase them. This distinction also clarifies the role of different margins: if long guns are less lethal than handguns, both within-category substitution and movement to the outside option reduce homicides, whereas if lethality differences are small, the extensive margin drives the results.

Averted Deaths Calculation. A policy that changes the household handgun ownership share from $G_{HG,0}$ to G'_{HG} and long-gun-only ownership share from $G_{LG,0}$ to G'_{LG} averts the following proportion of deaths:

$$\text{Averted deaths} = \left[H_{HG,0} - H_{HG,0} \cdot \left(\frac{G'_{HG}}{G_{HG,0}}\right)^{\varepsilon_{HG}}\right] - \left[H_{LG,0} \cdot \left(\frac{G'_{LG}}{G_{LG,0}}\right)^{\varepsilon_{LG}} - H_{LG,0}\right] \quad (21)$$

[4]Further evidence that most firearms are not purchased with criminal intent is that the median time from retail sale to recovery at a crime scene is approximately nine years (Bureau of Alcohol, Tobacco, Firearms, and Explosives).

The first bracket is deaths averted due to a reduction in handgun ownership; the second is increased *mortality* from greater long gun ownership (if $\varepsilon_{LG} > 0$).

5.4 Optimal Stopping Rule

We characterize the optimal policy by comparing marginal costs and benefits. Ordering individuals by their willingness to accept, the marginal cost of inducing the $(n+1)$th person to switch is their WTA. The marginal benefit is the value of averted deaths; we use a value of statistical life (VSL) of $13.1 million, consistent with current federal regulatory guidance (U.S. Department of Health and Human Services, 2024). The optimal subsidy equates the marginal cost and the marginal benefit:

$$\text{WTA}_{n^*} = \text{VSL} \cdot \frac{\partial(\text{Averted deaths})}{\partial n}. \tag{22}$$

A Note on the Mechanism. Before proceeding, we note an important feature of the mechanism. We hypothesize that gun-naïve households effectively have unit demand for firearms; because these households enter the market to address a specific concern – typically personal safety – additional units provide limited incremental value for that purpose. The analogy to home security systems is apt: one suffices. Under this hypothesis, the policy does not create scope for arbitrage or repeated take-up. A household that accepts compensation to choose a long gun reveals that the long gun, combined with the subsidy payment, adequately satisfies their underlying motivation for entering the market.

5.5 Perfect Targeting

The planner offers each gun-naïve buyer their exact WTA to switch away from a handgun purchase. In this scenario, the total cost of inducing n people to switch is:

$$C_A(n) = \sum_{i=1}^{n} \text{WTA}_{(i)} \tag{23}$$

where $\text{WTA}_{(i)}$ is the ith order statistic (the ith lowest WTA). This represents a lower bound on the cost of achieving a given level of substitution.

Table 9 reports optimal policies under alternative assumptions about long gun homicide elasticities. In both cases, the optimal policy is cost-effective: the cost per averted death (\$5–7 million) falls below the \$13 million VSL. Under the more optimistic assumption, the optimal policy operates at a larger scale because increases in long gun ownership do not offset the mortality benefits of reducing handgun ownership. Figures 2 and 3 display the deaths averted as a function of the number of switchers under the optimistic and pessimistic elasticity assumptions, respectively.

Table 9: Perfect Targeting Results

	Optimistic ($\varepsilon_{LG} = 0$)	Pessimistic ($\varepsilon_{LG} = 0.2$)
% of HG buyers switching	54.5%	33.8%
Maximum subsidy paid	\$959	\$435
Averted deaths	275	68
Total program cost	\$1.45 billion	\$0.47 billion
Cost per life	\$5.3 million	\$6.9 million

5.6 Uniform Subsidy

If the planner cannot observe individual WTA, they must offer the same subsidy S to everyone who purchases a long gun or the outside option (e.g., a taser). Because both alternatives are equally subsidized, the policy does not induce substitution between long guns and the outside option—relative attractiveness remains unchanged.

Under a uniform subsidy S, three types of households receive the subsidy: switchers (households that switch away from handguns), inframarginal long-gun buyers, and inframarginal outside-option buyers. Accordingly, the total policy cost of the uniform subsidy is:

$$C_B(S) = S \cdot \left(N_{\text{switch}}(S) + N_{LG}^{\text{orig}} + N_{out}^{\text{orig}} \right). \tag{24}$$

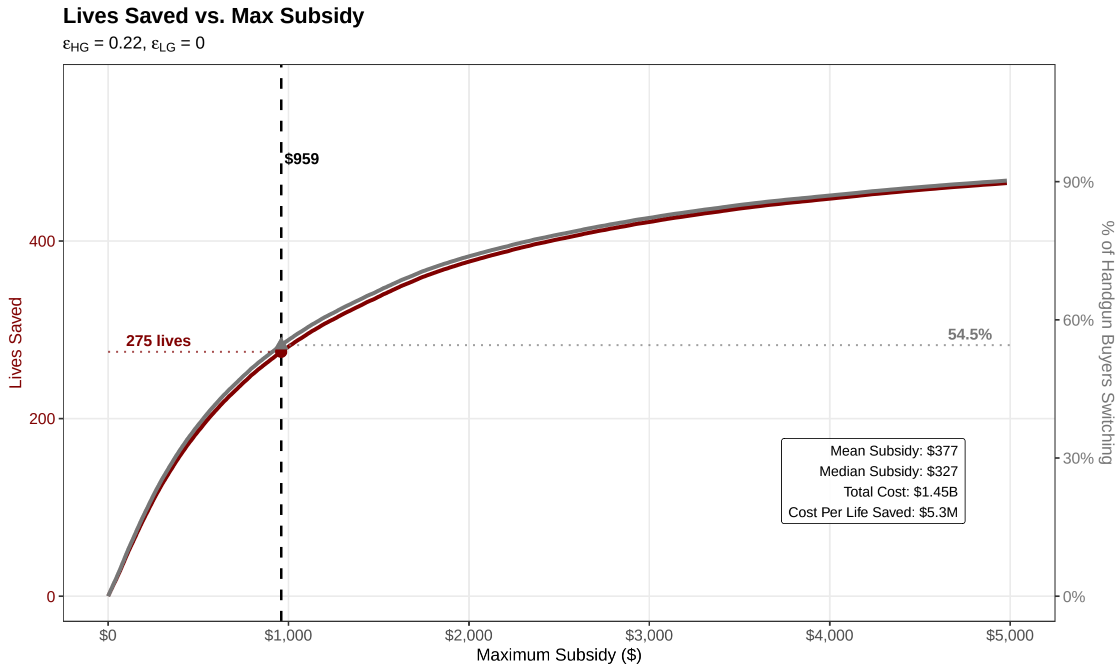

Figure 2: Averted Deaths Curve (Optimistic Scenario)

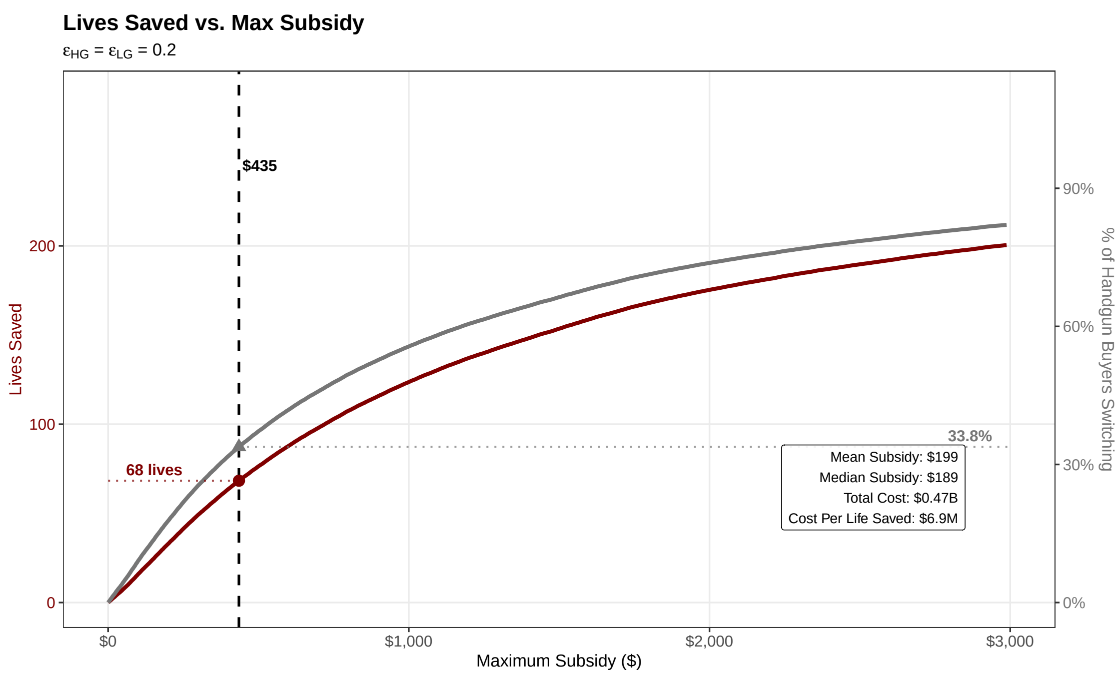

Figure 3: Averted Deaths Curve (Pessimistic Scenario)

This cost differs from the targeted subsidy for two reasons. First, switchers receive subsidy S regardless of their willingness to accept, leading to higher payments for low WTA customers. Second, inframarginal buyers receive the subsidy even though their behavior is unchanged; for these buyers, the subsidy is a windfall. Our predicted market shares imply that these windfall payments can be substantial. Among gun-naïve households, the predicted status quo long gun share is 28.7% and the predicted outside option share is 7.7%; any uniform subsidy therefore provides payments to these 36.4% of gun-naïve households who do not change their behavior.

Table 10 reports the optimal uniform subsidy under optimistic elasticity assumptions. The optimal uniform subsidy is much lower than the optimistic optimal targeted subsidy ($229 vs. $959), inducing fewer handgun buyers to switch to another alternative (20.1% vs. 54.5%). Total cost is comparable ($1.25B vs. $1.45B), but averted deaths are much lower (100 vs. 275). Nonetheless, the cost per life of $12.5 million is below the VSL threshold, implying that the policy is cost-effective.

Table 10: Scenario B: Uniform Subsidy Results (Optimistic)

Optimal uniform subsidy	$229
% of HG buyers switching	20.1%
Averted deaths (annual)	100
Total program cost	$1.25 billion
Cost per life	$12.5 million

Table 11 shows where the money goes. Only 26% of program payments are to bona fide switchers. The remaining 74% are windfall payments to inframarginal buyers who would have chosen long guns or the outside option even absent a subsidy. This leakage arises from the inability to target transfers.

Turning now to uniform subsidies under a pessimistic elasticity assumption ($\varepsilon_{LG} = 0.2$), we obtain a striking result: no positive uniform subsidy passes cost-benefit analysis. The minimum marginal cost per life is approximately $17 million, well above the $13 million VSL threshold. This occurs because the benefit per switcher is lower (79% substitute to

Table 11: Cost Decomposition: Uniform Subsidy at $229

Component	Cost	% of Total
Payment to switchers	$0.32 billion	26%
Windfall to LG buyers	$0.73 billion	58%
Windfall to outside buyers	$0.19 billion	16%
Total	**$1.25 billion**	100%

long guns, which increase homicides under pessimistic assumptions), while windfall payments to inframarginal buyers are unchanged. The result is that marginal cost exceeds marginal benefit at every positive subsidy level.

Table 12: Comparison of Policy Scenarios

	Targeted Subsidies		Uniform Subsidies	
	$\varepsilon_{LG} = 0$	$\varepsilon_{LG} = 0.2$	$\varepsilon_{LG} = 0$	$\varepsilon_{LG} = 0.2$
Optimal subsidy	$959	$435	$229	N/A
% switching	54.5%	33.8%	20.1%	N/A
Averted deaths	275	68	100	N/A
Total cost	$1.45B	$0.47B	$1.25B	N/A
Cost per life	$5.3M	$6.9M	$12.5M	>$13M
Cost-effective?	**Yes**	**Yes**	**Yes**	**No**

5.7 Information Rents

Table 12 compares targeted and uniform subsidies. For a given share of switchers (e.g., 55%), uniform subsidies cost 5–6 times as much as targeted subsidies. This difference reflects information rents paid to inframarginal buyers under uniform subsidies. For example, the optimistic uniform subsidy averts 100 deaths for $1.25B, compared to 275 deaths for $1.45B under perfect targeting. These rents also reduce robustness; targeted subsidies are cost-effective under both elasticity assumptions, whereas uniform subsidies are cost-effective only under optimistic assumptions.

We illustrate these forces in two figures. Figure 4 shows uniform subsidies reach only the left tail of the WTA distribution, whereas perfect targeting extends further. Figure 5 plots marginal benefit and marginal cost curves for both scenarios. The marginal benefit curves (maroon for optimistic, gray for pessimistic) are identical across policies, as they depend only on mortality effects. Marginal cost curves differ sharply: under perfect targeting (Scenario A), marginal cost equals the WTA of the marginal switcher; under uniform subsidies (Scenario B), it includes windfall payments to all inframarginal buyers. The gap between the marginal cost curves represents the information rent—the cost of not observing individual preferences.

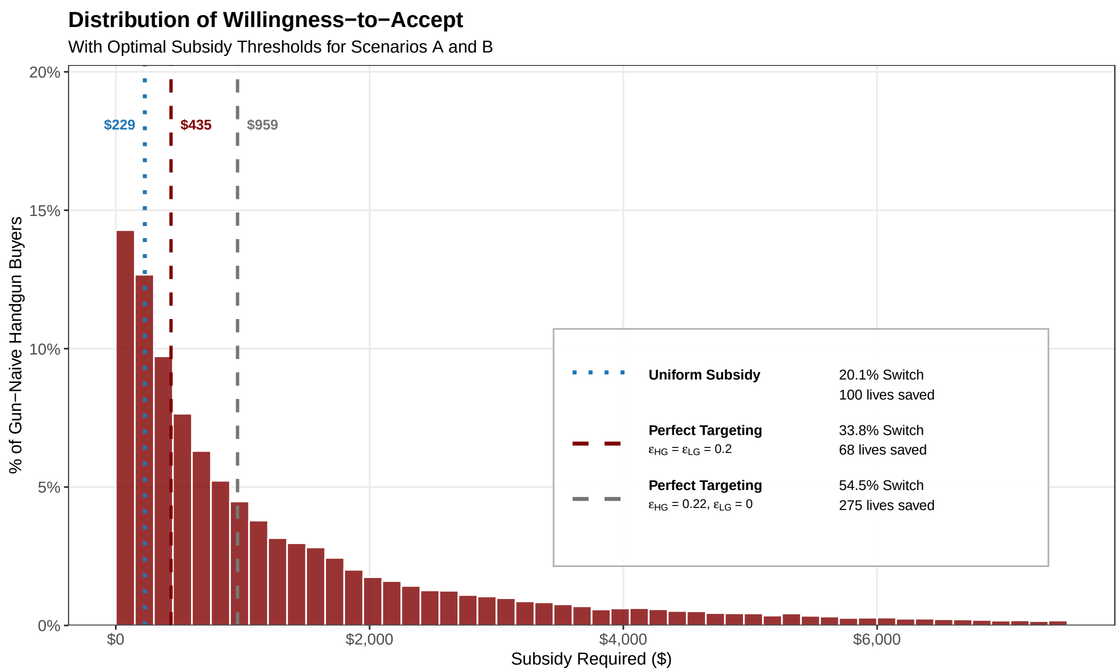

Figure 4: WTA Distribution with Optimal Subsidy Thresholds (Scenarios A and B)

5.8 Summary of Results

Our counterfactual analysis yields several key findings:

1. **Targeted compensation is cost-effective.** Under perfect targeting, compensating handgun buyers to choose alternatives costs \$5–7 million per averted death, below the VSL threshold under either elasticity assumption.

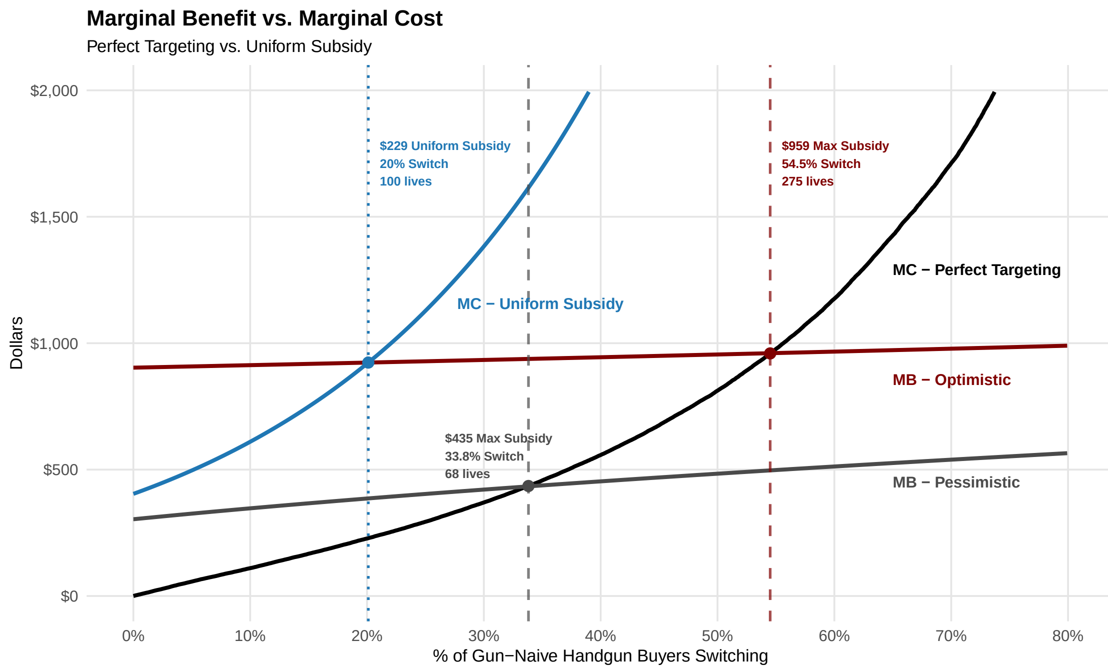

Figure 5: Marginal Benefit and Marginal Cost Curves (Scenarios A and B)

2. **Uniform subsidies are fragile.** Under optimistic assumptions, a $229 subsidy averts 100 deaths at $12.5M per life; under pessimistic assumptions, no subsidy is cost-effective.

3. **Information rents are large.** The gap between targeted and uniform subsidies reflects payments to inframarginal buyers. Perfect targeting averts 275 deaths for $1.45B; uniform subsidies avert only 100 deaths for $1.25B.

4. **Windfall payments dominate uniform subsidies.** Under uniform subsidies, 74% of expenditure goes to inframarginal buyers who would not have purchased handguns anyway.

These results suggest meaningful scope for compensation-based firearm policies, but they also highlight the importance of targeting and policy design. Even partial screening mechanisms may yield substantial improvements over uniform subsidies

6 Conclusion

Gun violence remains one of America's most intractable policy challenges. This paper studies an alternative approach: compensation rather than prohibition. We ask how much it would cost to pay potential handgun buyers to choose less lethal alternatives. Using a structural model of firearm demand, we estimate the distribution of willingness to accept for substitution away from handguns and computed the cost-effectiveness of various subsidy policies.

Our analysis yields three main findings. First, targeted compensation policies are highly cost-effective: under perfect targeting, optimal subsidies avert 68–275 deaths annually at a cost of $5–7 million per life, below the $13 million VSL threshold under either elasticity assumption. Second, uniform subsidies are more expensive and less robust; whether a cost-effective uniform subsidy exists depends on uncertain parameters, with no positive subsidy passing cost-benefit analysis under pessimistic assumptions. Third, the ability to target based on preferences is critical: at similar total expenditure, targeted subsidies avert more than twice as many deaths as uniform subsidies.

We do not propose a specific policy implementation. Implementing compensation-based firearm regulation would require addressing practical challenges, including verification, administration, political feasibility, and legal review. Our contribution is instead to quantify the *scope* for such approaches: the results indicate that compensation-based policies could reduce harm while making gun owners weakly better off.

More broadly, our analysis provides benchmarks for incentive-based policy designs. A substantial share of gun-naïve households would switch to long guns or the outside option for relatively modest compensation. This suggests that many prospective buyers may not be strongly attached to handguns per se. This pattern is consistent with demand driven by perceived safety benefits, which could potentially be provided by alternative products or services, such as non-lethal firearms, home security systems, or community safety programs.[5]

[5]Non-lethal firearms such as the Byrna, a CO2-powered launcher that fires pepper balls and other non-

While we do not estimate willingness to accept for such alternatives, the demand-side margins we document suggest they may be promising.

Several limitations point to directions for future research. Our analysis relies on stated choice data, which may not fully reflect realized behavior; validation using revealed preference data would strengthen the findings. We focus on partial equilibrium effects, abstracting from general equilibrium responses (e.g., secondary markets, supply-side pricing, criminal acquisition). Implementation also raises issues around arbitrage and fraud that we do not address. Finally, intermediate policy designs between perfect targeting and uniform subsidies—partial screening, menu-based approaches, conditional subsidies—might achieve better outcomes and merit further study.

lethal projectiles, offer a form factor similar to handguns while posing far lower mortality risk. Stantcheva et al. (2025) document potential consumer interest in such alternatives.

References

Armona, Luis and Adam Rosenberg, "Second-Best Amendment: Market Power and Tax Design in the Firearms Industry," 2025. Working Paper.

Blocher, Joseph and Darrell A.H. Miller, *The Positive Second Amendment: Rights, Regulation, and the Future of Heller*, Cambridge, UK: Cambridge University Press, 2018.

Braga, Anthony A. and Philip J. Cook, "The Association of Firearm Caliber with Likelihood of Death from Gunshot Injury in Criminal Assaults," *JAMA Network Open*, 2018, *1* (3), e180833.

Bureau of Justice Statistics, "Female Murder Victims and Victim-Offender Relationship, 2021," U.S. Department of Justice 2022.

Carlsson, Fredrik and Peter Martinsson, "Do Hypothetical and Actual Marginal Willingness to Pay Differ in Choice Experiments?," *Journal of Environmental Economics and Management*, 2001, *41* (2), 179–192.

Cook, Philip J. and Jens Ludwig, *Gun Violence: The Real Costs*, New York: Oxford University Press, 2000.

_ and _ , "The Social Costs of Gun Ownership," *Journal of Public Economics*, 2006, *90* (1-2), 379–391.

_ , _ , and Adam M. Samaha, "Gun Control after Heller: Threats and Sideshows from a Social Welfare Perspective," *UCLA Law Review*, 2013, *60*, 1104–1166.

Duggan, Mark, "More Guns, More Crime," *Journal of Political Economy*, 2001, *109* (5), 1086–1114.

Evans, William N., Craig Garthwaite, and Timothy J. Moore, "Guns and Violence: The Enduring Impact of Crack Cocaine Markets on Young Black Males," *Journal of Public Economics*, 2022, *206*, 104581.

Gallup, "What Percentage of Americans Own Guns?," Gallup Poll 2021.

Goss, Kristin A., *Disarmed: The Missing Movement for Gun Control in America*, Princeton, NJ: Princeton University Press, 2006.

Goulder, Lawrence H. and Ian W.H. Parry, "Instrument Choice in Environmental Policy," *Review of Environmental Economics and Policy*, 2008, *2* (2), 152–174.

Lacombe, Matthew J., *Firepower: How the NRA Turned Gun Owners into a Political Force*, Princeton, NJ: Princeton University Press, 2021.

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

Olson, Mancur, *The Logic of Collective Action: Public Goods and the Theory of Groups*, Cambridge, MA: Harvard University Press, 1965.

Rosenberg, Adam, "Regulating Firearm Markets: Evidence from California," 2025. Job Market Paper.

Rossi, Peter E., Greg M. Allenby, and Robert McCulloch, *Bayesian Statistics and Marketing*, Chichester, UK: John Wiley & Sons, 2005.

Stantcheva, Stefanie et al., "The Universal Pursuit of Safety and the Demand for Guns," 2025. Working Paper, Social Economics Lab.

U.S. Department of Health and Human Services, "Guidelines for Regulatory Impact Analysis," Office of the Assistant Secretary for Planning and Evaluation 2024. Value of Statistical Life: $13.1 million (2024 dollars).

A Estimation Details

This appendix provides technical details on the estimation of the search cost demand model. The exposition follows Moshary et al. (2025), which develops the full model and estimation procedure.

Model Setup. Consumers know their category-level preferences β_i' but must incur search cost γ_i to learn their idiosyncratic match ε_{ijt} for all models in a category. We impose a sign restriction so that $\gamma_i \geq 0$ for all consumers.

Inclusive Value. The inclusive value for category l for individual i is:

$$IV_{il} = \ln\left[\sum_{k \in l} \exp(X_k' \beta_i - \alpha_i \cdot \bar{p}_k)\right] \tag{25}$$

where $\bar{p}_k$ is the average price in category k.

Consideration Set Constraints. With four firearm categories, each consumer chooses one of the following consideration sets: (1) most preferred category only, (2) most and second-most preferred, (3) all but least preferred, or (4) all categories. The model implies the following revealed preference constraints for consideration set l_i:

$$IV_{il} \geq IV_{i,l+1} - \exp(\gamma_i) \tag{26}$$

$$IV_{il} - \exp(\gamma_i) \geq IV_{i,l-1} \tag{27}$$

$$\min_{c \in l_i} IV_{ic} \geq \max_{c \notin l_i} IV_{ic} \tag{28}$$

The first two inequalities ensure that the respondent does weakly worse by considering one more or fewer category. The third ensures that the worst considered category is weakly preferred to the best non-considered category.

Modified Likelihood. The likelihood is modified to include an indicator that the consideration set constraints hold:

$$P(y_{it}|\theta_i) = s_{ijt} \cdot P(\mathcal{C}_t|l_i) \cdot \mathbf{1}\{l_i|\theta_i\} \tag{29}$$

where s_{ijt} is the standard logit probability, $P(\mathcal{C}_t|l_i)$ is the probability of observing choice set $\mathcal{C}_t$ given stated consideration set l_i, and $\mathbf{1}\{l_i|\theta_i\}$ indicates that the consideration set constraints are satisfied.

Identification. To estimate the distribution of search costs, the conjoint includes choice tasks where respondents evaluate firearms from categories outside their stated consideration set. For each non-considered category (of which there may be up to three), one of the seven choice tasks is randomly selected to feature an alternative from that non-considered category. In 13.3% of these tasks, respondents choose the non-considered alternative, providing the identifying variation.

Hierarchical Prior. The search cost model uses the following priors:

$$V_\theta \sim IW(\nu, V) \tag{30}$$

$$\text{vec}(\Delta)|V_\theta \sim N(\text{vec}(\bar{\Delta}), V_\theta \otimes 100 \cdot I) \tag{31}$$

where $u_i \sim MVN(0, V_\theta)$.

MCMC Algorithm. The estimation proceeds as follows:

1) **Initialize.** For each respondent, run a separate logit based on the stated consideration set to obtain initial $\hat{\beta}_i$. For non-considered categories, draw from the truncated distribution satisfying the inequality constraints.

2) **Metropolis Step for θ.** For each respondent i, generate candidate draws $\tilde{\theta}_i \sim$

$MVN(\theta_{i(s)}, b^2 V_{\theta(s)})$ where $b = 0.66$ is a scaling parameter. Accept with probability:

$$a = \min\left\{1, \frac{P(Y|\tilde{\theta}_i)P(\tilde{\theta}_i|\Delta_{(s)}, V_{\theta(s)})}{P(Y|\theta_{i(s)})P(\theta_{i(s)}|\Delta_{(s)}, V_{\theta(s)})}\right\} \tag{32}$$

3) **Gibbs Step for Δ, V.** Draw hyperparameters from their conditional distributions:

$$\text{vec}(\Delta_{(s+1)})|V_{(s)}, \vec{\theta}_{(s+1)} \sim N\left((Z'Z + 0.01 \cdot I)^{-1}(Z'\theta_{(s+1)} + 0.01 \cdot \text{vec}(\bar{\Delta})), V_{(s)} \otimes (Z'Z + 0.01 \cdot I)^{-1}\right) \tag{33}$$

$$V_{(s+1)}|\Delta_{(s)}, \vec{\theta}_{(s+1)} \sim IW(\nu + n, V + S) \tag{34}$$

where $S = (\theta - Z\tilde{\Delta})'(\theta - Z\tilde{\Delta}) + 0.01 \cdot (\tilde{\Delta} - \bar{\Delta})'(\tilde{\Delta} - \bar{\Delta})$ and $\tilde{\Delta} = (Z'Z + 0.01 \cdot I)^{-1}(Z'\theta + 0.01 \cdot \bar{\Delta})$.

4) Return to step 2.

MCMC Details. The Markov chain is run for 500,000 iterations with a burn-in of 50,000 draws. We retain every 500th draw, yielding 900 posterior draws.

B Ex-Ante Willingness to Accept

This appendix presents an alternative computation of willingness to accept based on an ex-ante criterion in which the consumer re-optimizes their consideration set from scratch. The main text computes WTA under an ex-post "can't unsearch" interpretation: the consumer has already searched their status quo categories, and the subsidy can only add long gun categories to the consideration set. Handgun categories never leave. This section instead allows the consumer to re-optimize search given the subsidized inclusive values, potentially dropping handgun categories entirely.

Methodology. Under the ex-ante interpretation, the consumer re-optimizes their consideration set from scratch given the subsidized long gun inclusive values. Let $\mathcal{C}^*(D)$ denote the consumer's optimal consideration set when long gun inclusive values are scaled by the subsidy, i.e., $\widetilde{IV}_{ik} = IV_{ik} \cdot \exp(|\gamma_i| \cdot D)$ for each long gun category k, and the consideration set is determined by the greedy search algorithm with first-free search costs. We define the ex-ante WTA as the minimum D such that the consumer is diverted from purchasing a handgun through one of two paths:

- **Path A:** If no handgun category appears in $\mathcal{C}^*(D)$, the consumer never searches handguns and therefore never observes the product-level shocks ε_{ij} for handgun products. The consumer is diverted regardless of whether they ultimately choose a long gun or the outside option.
- **Path B:** If a handgun category remains in $\mathcal{C}^*(D)$, the consumer observes handgun product shocks and the subsidized long gun must beat the consumer's realized handgun utility with certainty:

$$\max_{j \in LG \cap \mathcal{C}^*(D)} (U_{ij} + |\gamma_i| \cdot D) > U_i^{HG}. \tag{35}$$

We compute D by binary search to \$0.01 precision.

The key distinction between the two approaches is what happens to the status quo consideration set. Under the ex-post ("can't unsearch") interpretation used in the main text, the consumer has already searched their status quo categories; the subsidy can add long gun categories but cannot remove handgun categories. Under the ex-ante ("re-optimize search") interpretation, the consumer re-optimizes from scratch, potentially dropping handgun categories entirely. When handguns drop from the consideration set (Path A), the consumer is diverted without needing to beat a specific handgun option. When handguns remain (Path B), the criterion is identical to the main text.

Results. Figure 6 plots the ex-ante WTA distribution and Figure 7 shows lives saved under the optimistic elasticity specification. The ex-ante WTA distribution is shifted modestly to the left relative to the main results. Table 13 reports the Scenario A results.

Table 13: Scenario A: Perfect Targeting (Ex-Ante WTA)

	Optimistic ($\varepsilon_{HG} = 0.22$, $\varepsilon_{LG} = 0$)	Pessimistic ($\varepsilon_{HG} = \varepsilon_{LG} = 0.2$)
% of HG buyers switching	56.5%	35.4%
Maximum subsidy	\$962	\$433
Averted deaths	286	70
Total program cost	\$1.48 billion	\$0.49 billion
Cost per life	\$5.2 million	\$6.9 million

Under perfect targeting, the ex-ante WTA yields modestly more averted deaths (286 vs. 275 under the optimistic scenario) at similar cost per life. The mean ex-ante WTA is \$2,055 (vs. \$2,111 for ex-post), the median is \$758 (vs. \$812), and 80.6% of switchers go to long guns (vs. 78.9%). The maximum subsidy at the optimal is nearly identical (\$962 vs. \$959), reflecting the fact that the optimal stopping rule equates marginal cost to VSL regardless of the overall WTA level; the ex-ante distribution simply reaches this threshold at a higher quantile (56.5% vs. 54.5%).

Under uniform subsidies (Scenario B), the optimistic optimal subsidy is \$221, inducing 20.9% of handgun buyers to switch at a cost of \$11.8 million per life saved. As in the main text, no uniform subsidy is cost-effective under the pessimistic elasticity. Figure 8 displays the marginal benefit and marginal cost curves.

Discussion. The ex-ante and ex-post WTA distributions are highly correlated (correlation = 0.93), and the qualitative conclusions of the paper are unchanged under the alternative criterion. Both interpretations are defensible from different modeling perspectives. The ex-post interpretation corresponds to a setting where the consumer has already incurred the cost of searching their status quo categories and cannot undo that search; the subsidy can only expand the consideration set. The ex-ante interpretation corresponds to a setting where the consumer has not yet committed to any search and re-optimizes given the subsidized inclusive values, consistent with the search cost model's timing assumptions. Under the ex-ante criterion, consumers who would drop handguns from their consideration set are diverted at lower cost (Path A), while those who would retain handguns face the same requirement as in the main text (Path B). The fact that results are robust to this choice reinforces the paper's main conclusions about the feasibility and cost of compensation-based handgun substitution policies.

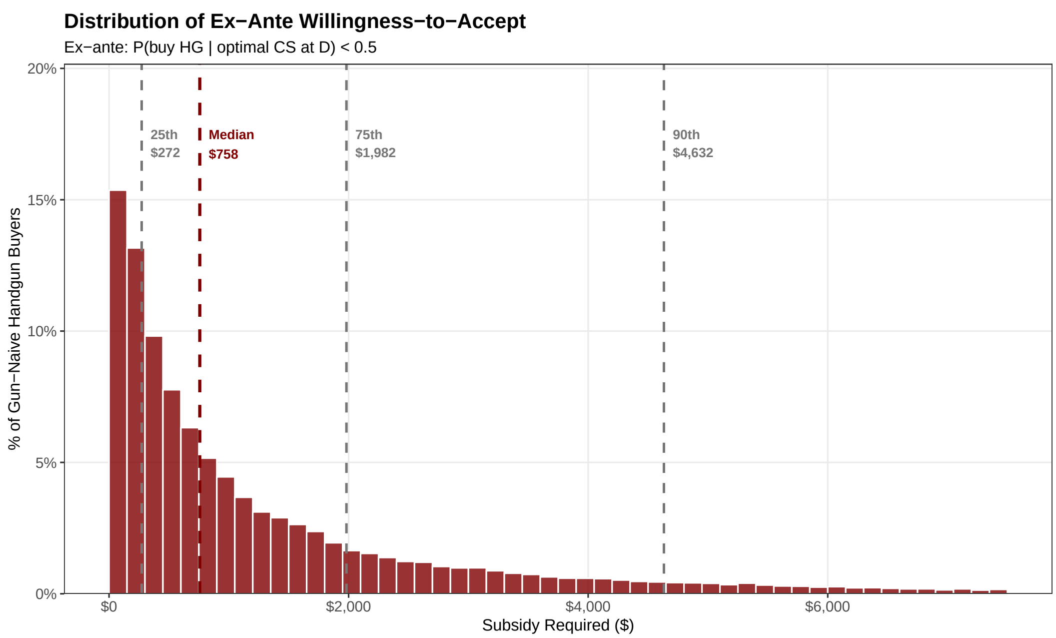

Figure 6: Distribution of ex-ante WTA among gun-naïve handgun buyers. Vertical lines mark the 25th, 50th (median), 75th, and 90th percentiles.

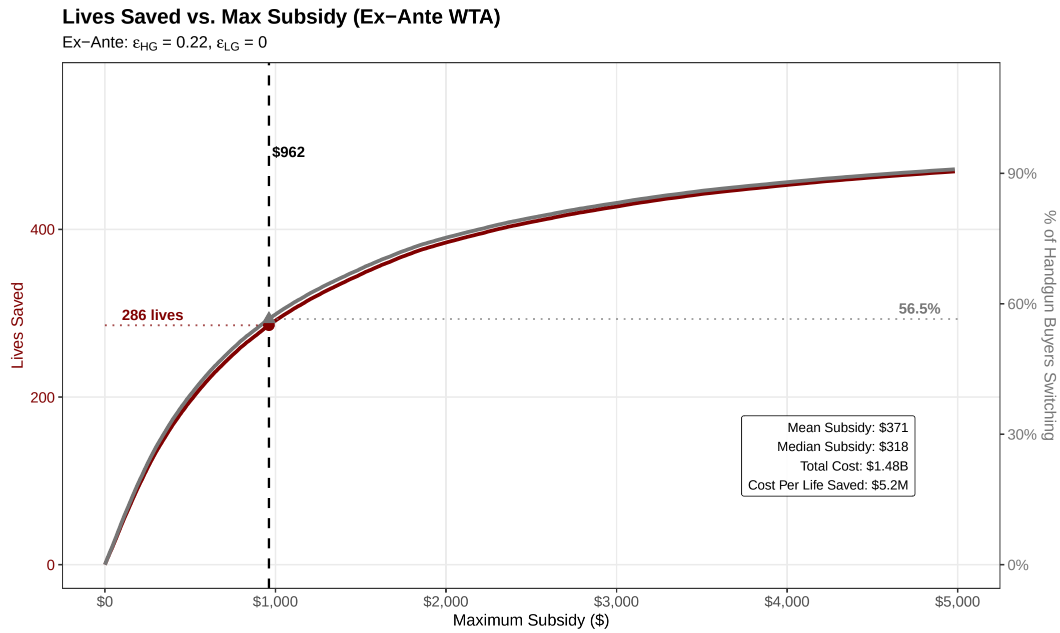

Figure 7: Lives saved as a function of the maximum subsidy under ex-ante WTA and the optimistic elasticity ($\varepsilon_{HG} = 0.22$, $\varepsilon_{LG} = 0$).

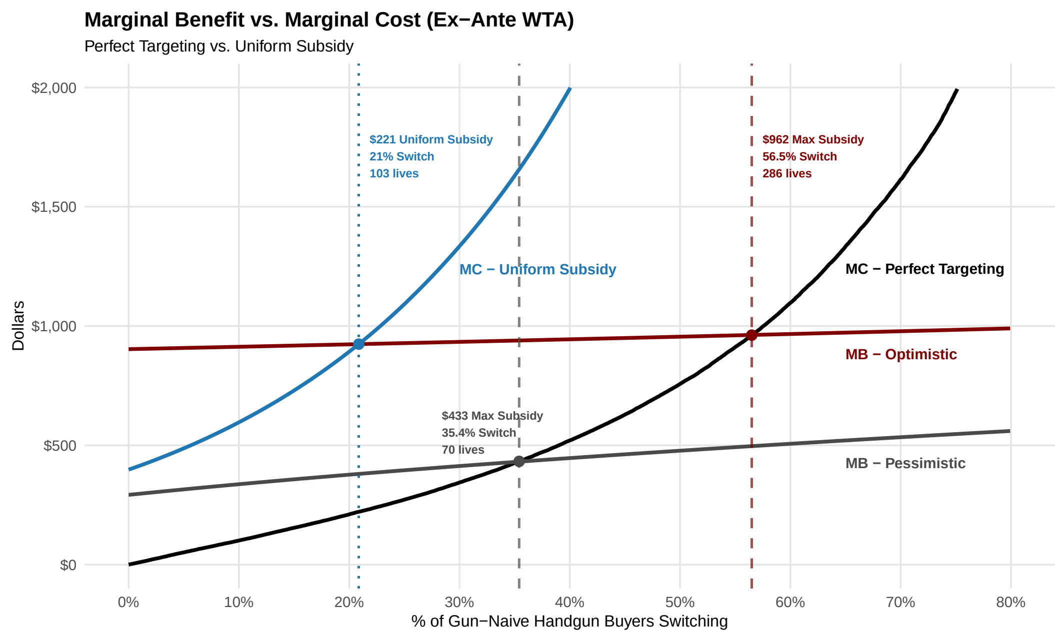

Figure 8: Marginal benefit vs. marginal cost under ex-ante WTA for both Scenario A (perfect targeting) and Scenario B (uniform subsidy).