

Cheap to Keep: Insurance Pricing and the Retirement of Hazardous Capital

Kaleb K. Javier*

[Click here for newest version](#)

October 2026

Abstract

In the midst of the clean energy transition, efficient coordination of hazardous capital retirement has become increasingly important. Much of the policy response operates on the adoption margin, and research on retiring capital already in place has focused largely on subsidies and mandates. This paper studies a different decommissioning instrument: pricing the cost of retaining hazardous capital through mandatory environmental insurance. I examine this question at U.S. gas stations, whose buried fuel tanks are the most commonly reported source of groundwater contamination in the country, and where federal law requires operators to carry pollution coverage but states differ in whether premiums are uniformly priced or risk-rated. I identify the effects of risk-based pricing using Texas's 1998 transition from a uniformly priced state insurance fund to risk-rated private insurance, which repriced 16,625 facilities. Using a panel of 5,587,134 tank-years and a difference-in-differences design that compares Texas tanks with tanks of the same construction and vintage in 23 control states, I find that risk-rated pricing increased annual tank closures by 1.3 percentage points, a 65 percent increase. The effect rises with tank age, from 0.4 points for the youngest quarter of tanks to 3.9 points for tanks 20 years or older at the reform. Small facilities disproportionately exited, while larger facilities replaced tanks and remained in operation. I then estimate a dynamic model of facility portfolio choice and entry to compare risk-based pricing with replacement subsidies and age-based retirement mandates. Among these instruments, risk-rated pricing comes closest to a first-best premium that prices both cleanup costs and health damages. Pricing mandatory environmental insurance is therefore a decommissioning policy, and one that outperforms the alternatives considered here.

*University of California, Berkeley, Department of Agricultural and Resource Economics. Email: kalebkja@berkeley.edu. Website: kalebkjavier.com. I thank Meredith Fowlie, Billy Pizer, James Saltee, Nano Barahona, Ben Handel, Matt Backus, and Cailin Slattery for their invaluable comments and feedback. I also thank the participants and discussants at Camp Resources XXXI and the AERE Summer Conference 2025 for helpful discussions. All remaining errors are my own.

1 Introduction

In the midst of the clean energy transition, efficient coordination of hazardous capital retirement has become increasingly important. Much of the economic research has focused on the adoption margin for the clean energy transition: subsidies for cleaner technologies, standards on new sources, and incentives for firms to replace older equipment with newer capital (e.g., [Anderson and Saltee 2011](#); [De Groote and Verboven 2019](#)). Much less attention has been given to the other side of the transition, getting hazardous capital that is already in place out. Policies designed to accelerate adoption of new technologies do not necessarily solve this exit problem: new capital can be subsidized while old capital remains grandfathered and continues operating. Research on capital retirement has consequently focused largely on subsidies that encourage owners to scrap old assets and mandates that require them to remove them (e.g., [Jacobsen and van Benthem 2015](#); [Barahona, Gallego, and Montero 2020](#); [Jacobsen et al. 2023](#)). This paper studies a different decommissioning instrument: risk-based pricing of mandatory pollution insurance.

The private value of continuing to operate an asset need not reflect the full social cost of keeping it in service. As capital ages, failure risk and expected environmental damages can rise even while the asset continues to generate revenue. An aging pipeline, fuel tank, oil well, or chemical storage system may therefore remain privately profitable after the point at which society would prefer it to be retired. Two failures separate the private value from the social cost. Third parties bear part of the damage from a release, and liability does not restore the owner's incentive when the damage exceeds the assets available to satisfy a claim ([Shavell 1986](#)). Mandatory insurance is the standard remedy, and its effect on retirement depends on how the coverage is priced. Because coverage is mandatory, a higher price cannot be avoided by leaving the insurance market, and the cost of continued operation can rise with the observable risk of the asset. This mechanism is not automatically efficient. The price must respond to the dimensions of risk that matter for expected harm, firms must change their decisions in response, and the price must not be so high that it retires capital whose social value still exceeds its risk. I therefore ask three questions. By how much does pricing risk change the retirement of hazardous capital? Does the response concentrate in the riskiest capital? And how does priced risk compare in welfare with replacement subsidies and retirement mandates?

I answer these questions at U.S. gas stations, whose underground fuel tanks are the most commonly reported source of groundwater contamination in the United States. The EPA counts about 583,000 confirmed releases from underground tanks since 1984, more than the 534,000 tanks now in service ([EPA 2026a](#)). Federal law requires every station to carry \$1 million of mandatory pollution insurance, but states differ in how they price that coverage. Most states historically charged tanks a uniform fee through a state insurance fund, while others relied on private insurers whose premiums vary with observable measures of tank risk, including age and construction. Texas provides a sharp change in those incentives. In 1998, the state closed its uniformly priced insurance fund to new claims and moved facilities into risk-rated private insurance, repricing 16,625 facilities. The timing of the transition was legislated in advance rather than chosen in response to facility behavior.

I identify the effect of risk-based pricing using a difference-in-differences design that compares retirement behavior in Texas before and after the reform with gas stations in 23 states that retained uniform pricing. Each Texas tank is compared with control-state tanks of the same construction, fuel, capacity, and age in the same calendar years, holding fixed the federal upgrade deadlines and the closure rules the control states enacted. The identifying variation therefore comes from the change in insurance pricing rather than from differences in the age or construction of the tanks being compared. I build the data from the universe of 402,673 registered tank facilities in 24 states, premium schedules collected by hand from insurer rate filings, and 12,168 claims obtained from 6 state funds. The estimates use a panel of 5,587,786 tank-years from 1990 to 2018. I then estimate a dynamic model of facility portfolio choice, in which owners maintain, downsize, replace, or exit and new facilities enter, to compare risk-based pricing with other retirement policies.

Risk-based pricing substantially increases retirement. The annual tank closure rate rises by 0.7 percentage points, 29 percent of the closure rate of comparable control-state tanks. The timing of the response is consistent with the reform: pre-reform coefficients are small, the series breaks at the transition, and the elevated retirement rate persists through the twenty post-reform years of the sample. The reform retired tanks rather than moving forward closures that would have happened anyway. By 2018 it had closed about 4,600 more of the 48,760 Texas tanks in service at the reform than would otherwise have closed, 9.5 percent of that stock and 29 percent more closures than it would otherwise have seen. A change in the price of mandatory coverage therefore moves the retirement margin. The reform itself changed no technical standard a tank must meet. The response is concentrated in the capital facing the largest increase in the price of continued operation. The effect rises with the age of the tank at the reform, from 0.2 percentage points for tanks 9 years old or younger to 3.6 percentage points for tanks 22 years or older, 84 percent of their control-state closure rate. A level shift in the cost of operating a station would tend to affect tanks more uniformly. The response instead follows the risk gradient embedded in the insurance price. A retirement policy is valuable when it causes the assets with the highest social cost of continued operation to leave. The reform moves the oldest tanks most, and those are the tanks the filed schedules charge most. The reform also changes how facilities adjust their portfolios. Exit rises by 0.4 percentage points a year, 23 percent of the control-state exit rate, and capacity cuts rise by a similar share. The tanks Texas closed were also cleaner. Releases discovered when a Texas tank was closed fell by 0.55 per 100 facility-years, nearly the full control-state rate.

The reduced-form evidence establishes that risk-based pricing changes retirement, but it does not by itself determine whether the resulting allocation is efficient or how the instrument compares with alternative ways of inducing exit. I therefore use the estimated dynamic model to evaluate three decommissioning instruments: risk-based pricing, replacement subsidies, and age-based retirement mandates. The model combines the observed price schedules with facility revenue, expected leak risk, and adjustment costs to characterize how operators choose whether to maintain, downsize, replace, or exit their tanks, and how many new facilities open. I measure each instrument against a first-best premium that prices both the cleanup cost and the health damage of a leak, counting incumbent facilities and entrants, by the share of the first best's gain over the observed contracts that it achieves. Extending risk-rated pricing to the states that kept a flat-fee fund achieves 45.8 percent of that gain.

The other instruments leave net welfare below the observed contracts. Returning Texas to a flat fee loses an amount equal to 8.3 percent of the gain, a replacement subsidy 2.6 percent, and a mandate requiring single-walled tanks to be removed at age 20 18.7 percent.

The paper makes four contributions. First, it provides facility-level evidence on how repricing environmental liability changes capital retirement, adding to research on liability and hazardous capital (Shavell 1986; Ringleb and Wiggins 1990; Yin, Kunreuther, and White 2007; Yin, Kunreuther, and White 2011; Yin, Pfaff, and Kunreuther 2011; Muehlenbachs 2015; Boomhower 2019; Marcus 2021; Boomhower et al. 2024). I observe a jurisdiction move its existing capital stock from a uniform fee to risk-rated pricing, and I measure which capital leaves. Second, it extends research on cost-based pricing in insurance markets (Einav, Finkelstein, and Cullen 2010; Bundorf, Levin, and Mahoney 2012; Hendren 2013; Einav, Finkelstein, and Mahoney 2021). That research studies how prices change who holds coverage. When coverage is compulsory, a cost-based premium changes which capital stays in service and acts as an imperfect corrective price. It charges owners for the expected insured loss and not for health damages. Both come from the same leaks, and a price on the first therefore reduces the second.

Third, the paper compares priced risk with scrappage subsidies and age mandates on a common welfare footing. Existing research on retiring durable capital has emphasized scrappage subsidies, retirement mandates, and vintage-differentiated regulation (Gruenspecht 1982; Adda and Cooper 2000; Stavins 2006; Mian and Sufi 2012; Jacobsen and van Benthem 2015; Barahona, Gallego, and Montero 2020; R. J. Davis, Holladay, and Sims 2022; Jacobsen et al. 2023; Deshmukh et al. 2023; L. Davis and Weber 2025). Fourth, it adds to research that uses dynamic models to study capital replacement and environmental regulation (Rust 1987; Ryan 2012; Kellogg 2014; Fowlie, Reguant, and Ryan 2016). Much of that work emphasizes sunk costs of entry and investment. I evaluate capital replacement under a liability cost that recurs in every year an asset stays in service.

The paper proceeds as follows. [Section 2](#) describes the tanks and the rules that govern them, the data, leak risk and its cost, the market for coverage, and the Texas reform. [Section 4](#) sets out the research design and reports the retirement response. [Section 5](#) develops and estimates the dynamic model, and [Section 6](#) compares the decommissioning policies. [Section 7](#) concludes.

2 Setting and Data

Institutional rules and technical standards for underground storage tanks (USTs) set the economic and regulatory environment for gasoline stations. Together, the regulatory framework, available technologies, and market incentives determine a facility’s risk profile and capital structure at the time of tank installation. These features lock in the choices operators can make over a tank’s life and shape the empirical analysis that follows.

2.1 Underground storage tanks and their regulation

Underground storage tanks (USTs) in the United States are constructed as either single-walled (SW) or double-walled (DW) systems, a choice made once at installation. The selected construction type determines the tank's long-term risk profile. A single-walled system consists of a single steel layer, making it susceptible to corrosion from both soil and fuel. In contrast, a double-walled system encloses the inner shell with a second, corrosion-protected layer. It includes a monitored interstitial space to detect breaches before fuel escapes into the environment. The leak risk for each tank is therefore established at installation and remains largely unchanged by subsequent operator behavior.

The choice between SW and DW construction entails a stark trade-off between upfront cost and risk mitigation, with no corresponding difference in revenue or operating capacity. In the late 1980s, double-walled systems cost approximately 70 percent more than single-walled alternatives. For a typical retail gasoline station operating three or four tanks, the incremental capital outlay for DW construction amounted to \$120,000 to \$160,000 in 2023 dollars. Crucially, the additional expenditure on DW systems did not increase storage, throughput, or revenue; its sole effect was to reduce the probability of fuel leakage. As expected, most operators selected the less costly SW option, reflecting the lack of immediate economic incentives to internalize future leak risk.

Once installed, a UST's construction type does not change during the tank's life. Operators can update piping or monitoring equipment, but converting from single-walled to double-walled construction requires full excavation and tank replacement, which is rarely done. Two institutional factors reinforce this outcome. Regulatory grandfathering exempts existing systems from stricter technical standards unless operators make major upgrades, so most have no incentive to replace legacy tanks. Excavation also risks exposing previously undetected contamination, which would trigger costly remediation. These constraints mean a facility's risk profile is set at installation and remains constant over time.

These institutional and technical constraints mean that capital in underground storage tanks is fixed at the point of installation. Firms have limited ability and little incentive to upgrade existing systems, facing both regulatory and financial barriers to replacement. As a result, adjustment occurs almost exclusively along the extensive margin, through tank closures or complete replacements, rather than through incremental investment in risk mitigation. This fixity of capital is central to the empirical setting and shapes both the design of the present study and the observed economic behavior in this market.

2.2 The Texas repricing

The fund failed on its own finances. Texas met the federal mandate through a fund that failed for reasons unrelated to the tanks it covered. The legislature created the Petroleum Storage Tank Remediation Fund in 1989 under the Texas Water Code. The fund reimbursed cleanups, charged a uniform \$50 per tank per year, and drew the balance from gasoline delivery fees. Its obligations outgrew that revenue. Actuarial review found the fund structurally unsustainable by 1995, and companion legislation that session authorized transferring up to \$120 million from general revenue to keep it paying claims ([House Research Organization 1995](#)).

The legislature closed it with three and a half years of notice. The legislature closed the fund rather than repair it. House Bill 2587 barred reimbursement of any release reported after December 22, 1998, making December 23 the first day under private coverage ([Texas Legislature 1995](#)). The sunset passed in 1995, which gave operators three and a half years of notice. 16,625 facilities moved from a uniform fee to risk-rated premiums on that date, in a single step.

The federal upgrade deadline and the Texas repricing fall on the same date, because the legislature set the fund’s closing date to match it. That deadline applied in all fifty states, and facilities in the comparison states met it while keeping their uniform fee. Its burden also fell unevenly across tanks. Release detection phased in by install cohort between 1989 and 1993, and tanks installed after 1988 owed nothing at all. Marcus ([2021](#)) uses the deadline itself as a treatment, to measure the health effects of leak exposure. This paper uses the repricing. [Section 4.1](#) states how the empirical design controls for the federal deadline.

Two features of that sequence carry the identification. The fund closed because of its own finances. Neither the decision nor its timing responded to the condition of any particular facility’s tanks. The fund also never repriced. It stopped writing entirely, and a private market already rating the same product in never-fund states absorbed its policyholders. The facilities in this study met a change in the price of risk that originated outside them. [Section 4.1](#) builds the comparison that measures their response.

2.3 Data

The analysis estimates empirical models on a panel of underground storage tank (UST) facilities in 25 states. The sample is restricted to these states because only they report tank installation date, closure date, and wall type consistently, three fields necessary to define event time and risk class. This restriction ensures comparability across facilities and a consistent estimation window. The final sample includes all facilities in these states that held at least one active tank at the December 23, 1998 reform; tanks installed later entered a risk-rated market and do not contribute before-and-after variation. Appendix [Section A.2](#) details the sample construction.

The analysis draws on four linked datasets, each matched at the facility level and used for specific empirical tasks:

UST Administrative Tank and Facility Universe: I use the EPA’s UST Finder database, which aggregates state registry records into a national schema ([EPA 2021](#)), to construct the universe of regulated tanks and facilities. This dataset provides the unit of analysis, facility-level tank portfolios, and all traits used to define risk classes (installation date, closure date, wall construction, fuel, capacity, location, and operator identity). Where coverage is incomplete, I supplement with state agency registry files. These data form the backbone of the estimation sample and supply the covariates for all reduced-form and structural models.

Rate Filings and Rating Manuals: I collected the complete rate filing history for the six largest carriers in the Texas market (2000–2025) from the Texas and Iowa insurance regulator archives. I reimplemented each carrier’s filed rating manual as an executable engine. These

data let me assign a risk-rated premium to every insured facility each year (both observed and counterfactual), anchoring all premium calculations and price comparisons in the analysis. Appendix [Section A.6](#) describes the engines and their coverage.

State Trust Fund Claims: I assembled the first facility-level, multi-state dataset of insurance claims on state remediation trust funds through records requests to six states: Colorado, Louisiana, New Mexico, Pennsylvania, Tennessee, and Utah. Each claim record reports the facility identifier, report date, and cleanup cost paid. Five of these states serve as controls in the closure analysis. I use these data to measure the distribution and expected value of cleanup costs after a release, which underpins all expected-cost calculations.

EPA Inspection and Compliance Measures: Every state submits biannual facility inspection and compliance statistics to the EPA, including the share of facilities in compliance with technical standards. I use these data to construct annual state-level enforcement measures for heterogeneity analysis and to compare institutional conditions across states. Appendix [Section B.11](#) details this construction.

3 Descriptive evidence

This section presents descriptive evidence on the risks associated with underground storage tanks (USTs) and the insurance market that covers those risks. The analysis first documents the incidence and distribution of leaks, detailing how risk varies by facility characteristics, tank construction, and regulatory environment. Facility-level data enable precise measurement of leak rates and cleanup costs across risk classes. The discussion then turns to the structure of the UST insurance market, tracking the number and specialization of active insurers, the formation and dispersion of premiums, and the alignment between observed risk factors and pricing. It also examines how premiums are set for facilities with features relevant to judgment-proofness and elevated risk.

Three central patterns emerge from the evidence presented below. First, leak risk varies substantially across facilities, driven primarily by tank construction type and age, with regulatory enforcement playing a secondary role. Second, the insurance market has few active insurers and shows significant variation in contract terms and premium schedules, only partially explained by observable risk. Third, certain high-risk or financially constrained facility types often face premiums that diverge from actuarial cost, indicating frictions in risk adjustment and institutional limits to full risk-based pricing.


3.1 Leak risk and its cost

A release is a cost to third parties before it is a cost to the operator. A release from a tank does not stay on the site. Gasoline reaches groundwater, and groundwater reaches drinking-water wells, neighboring parcels, and surface water. Underground tanks are the most commonly reported source of groundwater contamination in the United States ([EPA 2024](#)). The people who bear the damage are third parties. Marcus ([2021](#)) finds that infants born near a leaking tank site are more likely to be born preterm and at low birth weight, and the cleanup that follows a confirmed release is paid by the state fund or the insurer rather


than by the household exposed. The operator bears the deductible and the disruption. The risk this section measures is therefore a risk to society first, and the question the paper asks is whether the price the operator pays reflects it.

Figure 1: Leak risk and claims by tank construction

A. Effect of a single-walled tank on the annual release rate, by tank age



B. Claim rate and claim cost, by whether a facility holds a single-walled tank



Notes: Panel A: the effect of holding at least one single-walled tank on the annual probability that a facility reports a release, by three-year bin of the facility's average tank age, in percentage points. Each point is the single-walled coefficient plus its interaction with the age bin, from a linear probability model with state-by-year fixed effects, estimated on 3,146,158 facility-years in 22 states, Texas from 1990 to 1998 and the control states from 1990 onward. Bars are 95 percent confidence intervals from standard errors clustered by state. Wall type and age are measured in December. Panel B: claims to the state cleanup funds of Colorado, Louisiana, Pennsylvania, and Tennessee, in the years each state's claim records cover, from 1990 onward (Pennsylvania from 1994). Left: the percent of facility-years with at least one claim, with the number of facility-years with a claim and the number of facility-years printed above each bar. Right: the mean cost per claim in 2023 dollars, printed above each bar, with a 95 percent confidence interval from standard errors clustered by state (four clusters) and the median marked by a horizontal tick. In panels A and B the comparison group holds no tank recorded as single-walled and includes tanks with no recorded wall type.

Holding a single-walled tank raises the release rate as tanks age. I regress an indicator that

a facility reports a release in the year on the facility’s average-tank-age bin, an indicator that it holds at least one single-walled tank, and their interaction, with state-by-year fixed effects and standard errors clustered by state. Panel A of [Figure 1](#) plots the effect of holding a single-walled tank in each age bin. The effect is positive and its 95 percent confidence interval excludes zero in 5 of the nine age bins, peaking at +0.81 percentage points at ages 18-20. The effect grows with age through the late teens and is indistinguishable from zero in the first years of service. The comparison group holds no single-walled tank, and it includes tanks with no recorded wall type.

Release risk rises with age and is higher behind a single wall. The observed rate of first confirmed releases, tabulated by the facility’s average tank age in three-year bins and by wall construction, shows the pattern directly. Across the sample, 2.37 percent of facility-years record one. The rates show three facts. First, release rates rise with age. Among facilities with no single-walled tank the rate climbs from 1.01 percent at ages three to five to 1.99 percent past twenty-four. Second, single-walled facilities sit above facilities with no single-walled tank at every age. Their rate reaches 2.8 percent at ages eighteen to twenty, and the gap between the constructions is widest at ages 21-23, at 1.15 points. Third, the two ends of the profile are not age effects. Rates are highest in the first years of service, 4.7 percent for single-walled facilities and 4.0 percent for facilities with no single-walled tank, reflecting failures at installation rather than the ageing risk the rest of the paper prices, and the oldest single-walled bin falls as the leakiest tanks exit the sample.

Facilities holding a single-walled tank file claims more often. In the 4 states whose cleanup-fund claims link to the registry, Colorado, Louisiana, Pennsylvania, and Tennessee, a claim arrives in 1.17 percent of the facility-years of facilities holding a single-walled tank and 0.83 percent of those holding none (panel B of [Figure 1](#)). A claim costs \$353,099 at the mean in the first group and \$271,068 in the second, with medians of \$90,481 and \$77,020, but the state-clustered intervals on the two means overlap. Within state and year the difference in mean claim cost is +61,338 dollars ($p = 0.066$). The comparison group includes tanks with no recorded wall type.

A release costs \$106,000 at the median, and the tail is long. The claims data measure what a release costs. 12,168 claims paid by 6 state trust funds put a distribution on cleanup cost, plotted in [Figure 20](#) of [Section B.2](#). The median claim runs \$106,000 against a mean of \$348,000, and the gap between them reflects a right-skewed distribution in which a minority of releases accounts for most of the spending. The claims enter the analysis once, as the severity behind every expected-cost figure that follows.

Risk is predictable from the registry. To measure the risk of an individual facility rather than a cell, I estimate the annual probability that facility i in state s records its first confirmed release in year t :

$$\Pr(\text{Release}_{ist} = 1) = F(\alpha_t + \delta_s + \mathbf{X}'_{ist}\theta) \quad (1)$$

where F is the logistic function, the sample holds facility-years before any confirmed release, and the year and state effects α_t and δ_s absorb common shocks and fixed differences in

state reporting and enforcement. $\mathbf{X}_{ist}$ is the registry record. It holds the three-year bins of average tank age used above, an indicator for single-walled tanks, and their interactions, and the interactions let the age profile differ by wall type. It also holds the tank count, total capacity, fuel type, the roster’s composition by age, vintage, capacity, and fuel, and pairwise interactions among these traits. I estimate the equation as a lasso logit, a penalized logistic regression whose penalty is chosen by ten-fold cross-validation and whose selection shrinks uninformative terms to zero. The folds are grouped by facility, all of a facility’s years are held out together, and every prediction reported in this paper is out of fold. The model predicts releases well out of sample. It reaches an AUC of 0.81, the highest-risk decile of facilities contains 55.6 percent of all releases, and the predicted cell rates recover the observed age-by-wall profile cell by cell (Figure 16, in Section B.1). An insurer or a regulator observing only the registry can therefore assign each facility a risk before coverage is written.

At every age, a single-walled facility carries more expected cleanup cost than a facility with no single-walled tank. Combining the two measures prices the risk. A facility’s expected annual cleanup cost is its cell’s release rate times the mean cleanup cost of the claims in that cell, the same two objects the structural model prices in Section 5. I compute that cost in each age-by-wall cell, and the levels are measured rather than modelled: each is an observed rate times an observed mean claim. The two cells I quote differ in both age and wall. A facility with no single-walled tank whose tanks are three to five years old carries \$3,138 a year of expected cleanup cost. A single-walled facility whose tanks are twenty-four or older carries \$7,263, a factor of 2.31. Between those two cells, a single-walled facility carries more expected cleanup cost than a facility of the same age with no single-walled tank in every age bin.

The expected cleanup cost above puts a value on the two actions that reduce risk, closure and replacement. Consider a single-walled facility whose tanks are 29 years old, the median in 2015. Retiring five years early avoids \$24,511 of expected cleanup cost in present value, the annual cost above carried over five discounted years. Replacing with a new double-walled unit avoids \$23,036 over the new tank’s life. Replacement earns less than early retirement for a reason visible in the release rates above: a newly installed tank is not immediately safe, and those first years of the swap are the ones discounting weights most. Both figures are avoided expected cleanup cost rather than profit. Shedding tanks and operating on fewer is the third margin, and it removes risk only to the extent the tanks removed are the old ones.

3.2 The market for pollution coverage

In this subsection, I describe the private insurance market that replaced the State Fund in Texas. Federal rules require every tank owner to carry one million dollars of coverage per occurrence for cleanup and third-party liability. In all, 42 states created a State Fund to meet this requirement. A State Fund pays cleanup costs from motor fuel taxes and charges every tank the same annual fee, regardless of the tank’s age, wall construction, or release history. In 1998 the fee ranged from \$25 to \$500 across the 26 states that charged one, and Texas charged \$50 (ASTSWMO 1999). After the reform, Texas operators bought coverage from private insurers, which price each tank by the age and wall construction the state registry records. I use the Texas Commission on Environmental Quality’s financial-responsibility

records, which name each station’s insurer on July 1 of every year, and I price each station’s tanks with its insurer’s filed rating manual. I answer three questions. Who sells coverage? What does coverage cost, and how does its price vary across insurers and with a station’s risk? How large is the premium relative to the profit of the firms that buy it?


Few insurers sell coverage in Texas, and the market concentrated over time. In 2008, 18 insurers covered the 15,888 active gas stations in Texas. Of these insurers, 7 covered at least one percent of insured stations, and the four largest covered 78.3 percent. Tank Owners Members Insurance Company (TOMICS), a Texas insurer, held the largest share, 22.4 percent. Mid-Continent Casualty overtook TOMICS in 2010 and took over the renewal of TOMICS policies in January 2018. By 2020, Mid-Continent covered 57.8 percent of insured stations and the four largest insurers covered 84.2 percent. The number of insurers rose to 27, but only 10 of them covered at least one percent of insured stations. The HHI of insurer shares rose from 1,707 in 2008 to 3,662 in 2020. The Department of Justice and the Federal Trade Commission classify a market with an HHI above 1,800 as highly concentrated ([U.S. Department of Justice and Federal Trade Commission 2023](#)). Texas stayed below this threshold through 2015, at 1,581, and exceeded twice the threshold by 2020. Surplus-lines insurers, which are not licensed in Texas, covered 27.6 percent of insured stations in 2008 and 27.9 percent in 2020. Insurers also differ in the stations they cover. I classify a station as branded if its registry name contains a fuel brand, as in “Youngs Chevron”, and as unbranded otherwise.¹ In 2015 Mid-Continent covered 60.3 percent of priced unbranded stations and 43.7 percent of priced branded stations. Self-insurance grew. A firm may self-insure only if it has a tangible net worth of at least \$10 million and at least ten times its required coverage (40 CFR 280.95). The share of active stations that self-insured rose from 5.5 percent in 2008 to 10.0 percent in 2020. Total filed premiums for the stations I price fell from \$21.84 million in 2008 to \$10.91 million in 2020, in 2023 dollars.

Private coverage cost Texas stations several times the State Fund fee. I compare each station’s first filed premium with the fee it paid before the reform, \$50 per tank. The median station’s premium was 6.18 times its old fee, an increase of \$1,476 a year in 2023 dollars. Control stations kept their State Fund fee, and [Figure 2](#) compares the two groups before and after the reform. Before the reform, Texas stations paid less than control stations. In 1998 the median Texas station paid a State Fund fee of \$280 a year, 0.6 times the \$467 fee of the median control station. In 2015 the median Texas station’s filed premium was 3.6 times the \$386 fee of the median control station, and 36.9 percent of Texas stations had a premium above the fee that nine in ten control stations paid. These comparisons leave out the four control states whose State Fund charges no per-tank fee. Texas premiums varied widely and fell over time. The median premium was \$1,716 in 2008 and \$1,389 in 2015. In 2008 the premium at the ninetieth percentile, \$5,785, was 9 times the premium at the tenth percentile. Premiums fell as stations retired their oldest tanks ([Appendix C](#)). The insurer a station chooses explains little of this spread. In 2015 differences between insurers account for 4.9 percent of the variance in the premium per tank. The rest lies among tanks with the

¹I match facility names to the fuel brands of the Name Suggestion Index, the brand list OpenStreetMap maintains, among sites already classified as retail gas stations. I cannot classify 38.1 percent of Texas gas stations.

same insurer, which differ in age, wall construction, and the other traits the manuals rate.

Figure 2: The price of coverage per station, Texas and control states, 1998 and 2015



Notes: Kernel densities of the log annual amount per gas station, in 2023 dollars, on a log scale. In 1998 both groups paid a State Fund fee. In 2015 Texas stations faced a filed premium and control stations a State Fund fee. Texas in 1998: the State Fund fee of 50 dollars per tank times the station's active tanks, 16,670 stations. Texas in 2015: the filed premium of each priced station, the rating manual of its insurer applied to its tanks, 7,515 stations. Control states: the State Fund fee per tank times the station's active tanks, 43,555 stations in 1998 and 40,272 in 2015. Stations in the four control states whose State Fund charges no per-tank fee are not drawn, 32.7 percent of control stations with a recorded fee in 1998 and 31.6 percent in 2015. Vertical lines mark medians.

Premiums exceed the insurer's expected cost, and they sort stations by risk. I compare each station's filed premium with a benchmark for what the policy costs the insurer. The benchmark has two parts. The first is the station's expected cleanup cost. For every Texas station in every year, I predict the probability of a first confirmed release with the registry model of [Section 3.1](#), and I predict the cleanup cost of a release, net of the deductible, with a cost model fitted to the cleanup claims of that subsection on the same station traits. The expected cleanup cost is the product of the two, in 2023 dollars. The second is the minimum premium in the insurer's rating manual. Every manual sets a floor below which no policy's premium can fall, however few or safe a station's tanks are, and the floor averaged \$563 a year across insurers in 2015, in 2023 dollars. I treat this floor as the fixed cost of writing a policy. In 2015, 15.0 percent of Texas stations had a premium exactly at the floor. [Figure 3](#) plots the premium against this benchmark for every Texas station in 2008 and 2015. From 2008 to 2018 the premium averaged 1.30 times the benchmark. Texas stations recorded fewer first releases than the risk model predicts. Against the release rates Texas stations actually recorded, the premium averaged 2.42 times the benchmark. Insurers rank stations much as their risk does. The rank correlation between the premium and expected cost is 0.54 in 2008 and 0.62 in 2015. Above the floor, the premium rises about dollar for dollar with expected

cost for most stations. After I drop the top 1 percent of premiums and of expected costs in each year, each additional dollar of expected cleanup cost raised the premium by \$0.88 in 2008 and \$1.18 in 2015. These slopes measure expected cost at the risk model's release rates. Texas facilities recorded 0.269 first releases for each one the model predicts, so at recorded rates expected cost is that fraction as large. The floor leaves the safest stations facing premiums several times their expected cost. In 2015 the tenth of stations with the lowest expected cost faced an average premium of \$784 against \$175 of expected cost, while the riskiest tenth faced \$2,798 against \$2,106. A small group of the riskiest stations breaks this pattern. Their expected cost is very high, but their premiums are ordinary. With them included, each dollar of expected cost raised the premium by only \$0.21 in 2008 and \$0.51 in 2015. A State Fund ignores these differences. It charges an old single-walled facility the same fee as a young facility with no single-walled tank, although their expected cleanup costs differ by \$4,125 a year.

Figure 3: Filed Texas premiums against the insurer cost benchmark, 2008 and 2015



Notes: Each dot is one Texas gas station in the year shown: 7,427 stations in 2008 and 6,364 in 2015. The vertical axis is the filed premium, the rating manual of the facility's insurer applied to its tanks. The horizontal axis is the insurer cost benchmark: the facility's expected cleanup cost, a release rate times a fitted cleanup cost, plus the insurer's filed minimum policy premium. The dashed line marks a premium equal to the benchmark. The solid line is the least-squares fit of log premium on log benchmark in each year, over all of that year's stations. Both axes are on a log scale in 2023 dollars and show the middle 99 percent of values. The few stations outside that range are not drawn but enter the fit. Stations whose tanks average under three years of age are excluded. The log scales compress the largest premiums. In dollars, the correlation of the premium with the benchmark is 0.17 in 2008 and 0.26 in 2015.


Single-site owners hold a growing share of Texas stations. At the reform, 6,435 single-site

owners held 36.2 percent of active Texas gas stations, and 30 owners of more than fifty sites held 4,904 stations, 27.6 percent. From 1998 to 2018 the share of stations held by owners of two to fifty sites fell from 32.8 to 21.5 percent. Over the same period the single-site share rose to 44.4 percent and the share of owners of more than fifty sites rose to 28.9 percent. The rate per tank differs little by owner size. From 2015 to 2018 the mean filed premium per tank was \$568 for single-site owners and \$588 for owners of more than fifty sites, in 2023 dollars, and the mean filed premium per station was \$1,430 and \$1,544. The owner groups differ in the tanks they keep. [Figure 4](#) tracks the tanks of each owner group from 1995 to 2018, grouping owners by their size at the reform. Large chains retired single-walled tanks faster. At stations owned by chains of more than fifty sites, the single-walled share of active tanks fell from 86.0 percent in 1998 to 70.9 percent in 2018. At single-site stations it fell only from 89.8 to 84.7 percent. By 2018 the tanks at single-site stations averaged 32.1 years, against 28.7 years at chain stations. Realized release rates barely differ. From 2008 to 2018, single-site stations recorded 0.15 first releases per 100 station-years and chain stations 0.14. Single-site stations break the rules more often. From 2015 to 2018 the Texas Commission on Environmental Quality inspected both groups at about the same rate, 0.35 and 0.39 times per station-year, but 23.2 percent of routine compliance inspections at single-site stations found a violation, against 7.4 percent at chain stations.

The premium is a far larger burden for single-site owners than for chains. Relative to sales, the premium is small. In 2008 the average Texas station's filed premium equaled 0.046 percent of the average receipts of a Texas gas station in the Economic Census and 0.32 percent of the gross margin those receipts imply. Relative to profit, the premium is large for small owners. I compare each owner group's average premium with the profit of a typical firm of its type. For single-site owners, I use the mean profit per sole proprietorship return in gasoline stations that the Internal Revenue Service reports ([IRS 2008](#)). For owners of more than fifty sites, I use the operating profit of an average gasoline station that Census Bureau retail data imply ([U.S. Census Bureau 2024, 2008](#)). In 2008 the average single-site station's filed premium was \$1,475, 12.4 percent of single-store profit. The average station of a large owner had a premium of \$2,253, 1.3 percent of its operating profit. By 2015 these shares fell to 6.9 and 0.6 percent, as premiums fell and profits rose. Both benchmarks are national averages for each type of firm, not the accounts of Texas owners, and the premiums are in current dollars.

The reform replaced a flat fee with a price that rises with risk, and it hit Texas firms unevenly. Every Texas station faces a premium several times the fee a control station pays. Single-site owners carry the largest burden relative to profit, and they hold the older, single-walled tanks the premium prices highest. Chains face the same rate per tank, earn far more per station, and retired their single-walled tanks faster. The premium sends its strongest signal to the owners of the oldest tanks, even though a small group of the riskiest stations faces premiums far below its expected cost. The rest of the paper asks how operators with old tanks responded to this charge.

Figure 4: Single-walled share and tank age by owner size, 1995 to 2018



Notes: Texas gas stations active on 23 December 1998, grouped by the number of Texas sites their owner held on that date, and followed while active through 2018. Each point is 31 December of the year. Stations at the reform: 6,400 single-site, 3,264 with owners of 2 to 10 sites, 2,544 with owners of 11 to 50 sites, and 4,906 with owners of more than 50 sites. Before 1998 the series covers the same stations, so it contains only stations that survived to the reform. The left panel is the single-walled share of active tanks with a recorded wall type. The right panel is the mean age of active tanks. Ages are understated because 12.9 percent of Texas tanks carry the registry-start install date of 31 August 1987. The dashed line marks the reform.

4 The Closure Response to Risk-Rated Pricing

This section provides reduced-form evidence on how gas stations responded when Texas replaced its flat-fee fund with risk-rated private insurance. The unit of analysis is the tank. I define a tank's type as its wall construction, fuel, and capacity class, the traits the filed rate schedules price alongside age. The sample holds every tank at a gas station that was in service on December 23, 1998, the first day of private coverage, in Texas and the 23 control states whose funds remained in place. That stock contains 266,610 tanks at 78,483 facilities, 55,569 of them in Texas. I follow each tank from 1990 until it closes or the sample ends in 2018. The main outcome is permanent closure, the removal or closure in place of a tank as recorded in the state registry.

Texas and control-state stations were alike at the reform. The median station held 3 tanks in Texas and 3 in the control states, and the mean tank was 15.0 years old in Texas against 14.4 in the controls. In Texas 0.91 of tanks were single-walled against 0.89 in the controls, and 0.76 were installed before 1989 against 0.70. [Table 7](#) in [Section B.3](#) reports the full comparison.

Three facts emerge from the evidence below. First, Texas tanks closed more often than

comparable control-state tanks after the reform, and the gap persisted through the end of the sample. The reform therefore retired tanks that would otherwise have stayed in service rather than advancing closures that would have occurred anyway. Second, the response rose with the age of the tank at the reform, the trait the risk-rated schedules price most steeply. Third, stations responded by exiting and by cutting capacity, and the tanks Texas closed were consistently cleaner when inspected.

4.1 Changes in tank closure

I quantify the effect of the reform on tank retirement with an event-study design on the tank-year panel. I estimate

$$Y_{it} = \sum_{k \neq -1} \beta_k \cdot \text{TX}_i \cdot \mathbf{1}\{t - 1998 = k\} + \lambda' M_{st} + \alpha_{f(i)} + \delta_t + \eta_{a(i,t)m(i)} + \varepsilon_{it}, \quad (2)$$

where Y_{it} indicates that tank i is permanently closed in year t and TX_i indicates a Texas tank. Event time counts years from 1998, the year the reform took effect, and I omit $k = -1$. M_{st} holds the three federal deadline indicators, which switch on by install cohort, and 38 indicators for the closure mandates, upgrade mandates, and removal subsidies the control states enacted between 1991 and 2018, each on from its effective year. $\alpha_{f(i)}$ is a facility fixed effect and δ_t a year fixed effect. η_{am} is a fixed effect for each five-year age bin a of each tank type m . A Texas tank is therefore compared with control-state tanks of the same type at the same age. Standard errors are clustered by state.


Identification requires that, absent the reform, Texas tanks would have closed on the same path as control-state tanks of the same type and age. The legislature set the fund's closing date to the fund's revenue, and neither the decision nor its timing responded to the condition of any facility's tanks. The federal upgrade deadline of December 22, 1998 applied in every state, and the deadline indicators and the age-by-type effects absorb the response it induced in each cohort and construction. Closure mandates in control states enter through M_{st} . [Section B.4](#) discusses the remaining threats: state enforcement, anticipation of the 1995 sunset, and the double-wall standard for new tanks.

[Figure 5](#) displays the estimates of Equation 2. In the pre-period, the coefficients are small. They average 0.2 points and the largest is 0.4 in absolute value. The coefficient for 1998 is negative because Texas facilities removed fewer tanks at the federal deadline than control-state facilities did. After the reform, a Texas tank is more likely to close than a control-state tank of the same type and age in every year. The gap peaks at 1.7 points in 2003, 78 percent of the control-state closure rate that year, and averages 0.6 points from 2011 to 2018, 26 percent of the control rate over those years. The reform raised the closure rate of Texas tanks, and the increase lasted through the end of the sample.

Replacing the event-time terms with a single post-reform indicator gives

$$Y_{it} = \beta \cdot \text{TX}_i \cdot \text{Post}_t + \lambda' M_{st} + \alpha_{f(i)} + \delta_t + \eta_{a(i,t)m(i)} + \varepsilon_{it}, \quad (3)$$

Figure 5: Event study of tank closure, Texas relative to control states



Notes: Coefficients β_k from Equation 2, estimated on the tank-year panel. The outcome is an indicator that the tank is permanently closed in the year, and coefficients are in percentage points. 1997 is the omitted year. The sample is the 5,587,786 tank-years from 1990 to 2018 of the tanks at gas stations in service on December 23, 1998, in Texas and 23 control states. The regression includes facility and year fixed effects, a fixed effect for each five-year age bin of each tank type (wall construction, fuel, and capacity class), the three federal deadline indicators, and the control-state policy indicators. The dashed line marks the reform. Vertical bars are 95 percent confidence intervals from standard errors clustered by state.

with Post_t equal to one from 1999. A Texas tank closes 0.7 points more often each year after the reform (s.e. 0.2), 29 percent of the control-state closure rate of 2.4 percent over the same years.

4.1.1 More closures, not just earlier ones

A higher closure rate after the reform has two possible meanings. The reform could have retired tanks that would otherwise have stayed in service for many more years. Or it could have moved forward closures that would have happened a few years later anyway, leaving the total number of tanks retired unchanged. The two readings make different predictions for the share of the tanks in service at the reform that have closed by a given year. Under the first, the Texas share pulls ahead of what it would have been and stays ahead. Under the second, it pulls ahead and then falls back as the closures it moved forward stop arriving.


Measuring that share requires care because closure is permanent. A tank that closes leaves the stock, and yearly closure rates therefore cannot simply be added across years. I follow Deaner and Ku (2024) and estimate the comparison on the cumulative hazard, whose yearly

increments do add up. I group the tanks in service at the reform by type, install year, and state, and estimate

$$-\log(1 - h_{gt}) = \sum_{\tau \neq 1997} \theta_{\tau} \cdot \text{TX}_g \cdot \mathbf{1}\{t = \tau\} + \phi_g + \delta_t + u_{gt}, \quad (4)$$

where h_{gt} is the share of group g 's open tanks that close in year t , and ϕ_g and δ_t are group and year fixed effects. Observations are weighted by the number of tanks at risk, and standard errors are clustered by state. Each coefficient θ_{τ} is the reform's addition to the closure hazard of open Texas tanks in year τ .

Figure 6: Event study of the cumulative closure hazard, Texas relative to control states



Notes: Coefficients θ_{τ} from Equation 4, the effect of the reform on the yearly increment to the cumulative closure hazard, $-\log(1 - h)$, of Texas tanks relative to control-state tanks of the same type and install year, in percentage points. 1997 is the omitted year. The sample is the tanks at gas stations open at the end of 1998 in Texas and the 23 control states, grouped by wall construction, fuel, capacity class, install year, and state, and weighted by tanks at risk, with group and year fixed effects, 1990 to 2018. The dashed line marks the December 23, 1998 reform. Vertical bars are 95 percent confidence intervals from standard errors clustered by state. The sum of the post-reform coefficients through year k is ΔH_k in Table 9.

Figure 6 displays the estimates of Equation 4. In the pre-period, the coefficients are small. They average 0.3 points, and none differs significantly from zero. The 1998 coefficient is negative, as in Figure 5. Every coefficient after the reform is positive, between 0.1 and 1.6 points a year, and the largest fall in 2001 to 2003. The reform raised the closure hazard of Texas tanks in every year after it took effect.

Because the hazard increments add up, the coefficients give the number of tanks the reform closed directly. The coefficients from 1999 through 2003 sum to 5.1 points, the reform’s addition to the Texas cumulative hazard over its first five years. Applied to the Texas tanks still open, that addition closed about 2,100 more of the 48,760 Texas tanks in service at the reform than would otherwise have closed (s.e. 700).² By 2003, 8,100 Texas tanks had closed, against 5,900 without the reform, an increase of about a third. Through 2018 the coefficients sum to 15.4 points and the reform had closed 4,600 more tanks (s.e. 1,800). By 2018, 43 percent of the Texas tanks in service at the reform had closed, against 33 percent without the reform. The reform therefore retired an additional 9.5 percent of the 1998 Texas stock, 29 percent more closures than the stock would otherwise have seen.

If the reform had only moved closures forward, the later coefficients would turn negative as the advanced closures stopped arriving, and the running sum would fall back. No post-reform coefficient is negative, and the sum rises through the end of the sample. The reform retired tanks that would otherwise have stayed in service. Table 9 in Section B.6 reports the cumulative effects by install year, and Figure 21 the event study for each install cohort. The lasting effect is concentrated in tanks installed before 1985. For tanks installed from 1985 to 1988, the reform moved closures earlier without adding to them.

4.2 The role of risk pricing

If the response runs through the price of risk, the tanks whose premium rose most should close most. The filed schedules raise the premium most steeply with age (Section 3.2), while the fund charged every tank the same fee. I thus split the tanks in service at the reform into four groups by their age at the reform: 9 years or younger (77,223 tanks), 10 to 14 years (67,701), 15 to 21 years (59,278), and 22 years or older (62,408). I estimate

$$Y_{it} = \beta \text{TX}_i \text{Post}_t + \sum_{q=2}^4 \gamma_q \text{TX}_i \text{Post}_t \mathbf{1}\{q_i = q\} + \sum_{q=2}^4 \mu_q \text{TX}_i \mathbf{1}\{q_i = q\} \quad (5)$$


$$+ \lambda' M_{st} + \alpha_{f(i)} + \delta_t + \eta_{a(i,t)m(i)} + \varepsilon_{it},$$

where q_i is the tank’s age group, the youngest group is the reference, and all other variables and specification details are defined as in Equation 3. The terms μ_q allow Texas tanks of each age group to close at a different rate throughout the sample.

Results from Equation 5 are shown in Figure 7. Each point is the reform effect for one age group, and the dashed line is the pooled estimate from Equation 3. Tanks 9 years or younger at the reform close 0.2 points more often after the reform (s.e. 0.2), 20 percent of the closure rate of comparable control-state tanks, and the estimate does not differ significantly from zero. The effect rises with each group, to 1.1 points (s.e. 0.3) for tanks aged 10 to 14 and 1.8 points (s.e. 0.3) for tanks aged 15 to 21. Tanks 22 years or older close 3.6 points more often (s.e. 0.3), 84 percent of their control-state closure rate. These findings suggest that the

²The extra share closed after k years is $S_k (e^{\Delta H_k} - 1)$, where ΔH_k is the sum of the post-reform coefficients through year k and S_k is the share of the Texas stock still open. Multiplying by the stock gives the number of tanks.

Figure 7: Reform effect on annual tank closure by tank age at the reform



Notes: Each point is the estimated effect of the reform on the annual probability of tank closure for the tanks in one age group, $\beta + \gamma_q$ from Equation 5 and β for the youngest group, in percentage points. Age at the reform is 1998 less the install year, and the groups are 9 years or younger, 10 to 14, 15 to 21, and 22 or older. The dashed line is the pooled estimate from Equation 3. The group effects exceed the pooled estimate because old Texas tanks closed less often than old control-state tanks before the reform, a difference the terms μ_q absorb. Sample and controls follow Figure 5. Vertical bars are 95 percent confidence intervals from standard errors clustered by state. Table 10 reports the coefficients.

reform worked through the price of risk, with the largest response among the tanks whose premium rose most.

4.3 Margins of adjustment

A station facing a higher premium can close a tank and keep operating, replace it, consolidate its tanks, cut its capacity, or exit. I estimate Equation 3 at the facility level, one regression per margin, with each outcome an indicator for the facility-year. The facility specification mirrors the tank one. A facility's type is the type of its oldest tank crossed with its number of tanks at the reform, in four bins of one, two to three, four to six, and seven or more. The age-by-type effect becomes a fixed effect for each age bin of the oldest tank within each facility type, and the deadline and policy controls are unchanged. Table 1 reports the estimates.

The probability that a facility closes at least one tank rises 0.5 points (s.e. 0.2), 19 percent of the control-state rate of 2.7 percent. Two margins carry most of that response. Exit rises

Table 1: Facility-level estimates by portfolio margin

	(1) Any closure	(2) Exit	(3) Downsize	(4) Consolidate	(5) Reconfigure up	(6) Replace	(7) Capacity cut
Texas $\times$ Post	0.52 (0.20)	0.43 (0.17)	0.18 (0.07)	-0.11 (0.04)	-0.05 (0.03)	0.28 (0.09)	0.52 (0.18)
Control mean, post-reform (%)	2.70	1.87	0.46	0.11	0.09	0.47	2.44
Facility-years	1,882,224	1,882,224	1,882,224	1,882,224	1,882,224	1,882,224	1,882,224

Notes: Facility-level estimates of Equation 3, one regression per column, on the annual facility-year indicator named in each column head, in percentage points. The sample is the facility-years from 1990 to 2018 of gas stations holding at least one tank in service on December 23, 1998, in Texas and the 23 control states. Each regression includes facility and year fixed effects, a fixed effect for each five-year age bin of the facility’s oldest tank within each facility type (the oldest tank’s wall construction, fuel, and capacity class crossed with the facility’s number of tanks at the reform), the three federal deadline indicators, and the control-state policy indicators. Section A.3 defines each margin. The control mean is the post-reform control-state rate. Standard errors clustered by state are in parentheses.

0.4 points (s.e. 0.2), 23 percent of the control-state exit rate of 1.9 percent, and capacity cuts rise 0.5 points (s.e. 0.2), 21 percent of the control rate of 2.4 percent. Replacement rises 0.28 points (s.e. 0.09) and downsizing 0.18 points (s.e. 0.07), although Texas facilities replaced and downsized somewhat more than control facilities before the reform as well (Section B.8). Consolidation falls 0.11 points (s.e. 0.04).

The reform moved stations out of the market and shrank the ones that stayed. Wall construction cannot be changed on an installed tank, and a station seeking a lower risk class removes the tank rather than modifying it. Entry of new stations shows no perceptible response to the reform (Section B.10).


4.4 Confirmed releases

Finally, I ask how the reform changed the releases found at Texas facilities. The outcome is an indicator that a confirmed release is discovered at the facility in the year. Closing a tank triggers an inspection of the ground beneath it, and a release found there shows whether the closed tank had leaked. I therefore separate releases discovered within sixty days of a closure from those discovered outside a closure. I estimate Equation 2 on each outcome at the facility level, with the controls of Section 4.3.

Figure 8 displays the estimates. Panel (a) pools all confirmed releases. In the early 1990s, Texas facilities discovered more releases than control-state facilities, by up to 2.5 points in 1992, and the gap narrows through 1996. After the reform, the coefficients are close to zero in every year.

Panel (c) shows releases discovered at a closure. Before the reform, the coefficients are small and stable, between 0.1 and 0.3 points. They fall by 1.0 point in 1999 and keep falling, to 0.6 points below the control states by 2016 to 2018. On average after the reform, a Texas

Figure 8: Event study of confirmed release discovery, Texas relative to control states



Notes: Coefficients on the interaction of a Texas indicator with an indicator for each year relative to 1998, from Equation 2 estimated on the facility-year panel with the controls of Table 1. The outcome is an indicator that a confirmed release is discovered at the facility in the year. Panel (a) pools every confirmed release, panel (b) keeps releases discovered outside a closure, and panel (c) keeps releases reported within the sixty days following a closure. Coefficients are in percentage points. 1997 is the omitted year. The sample is the facility-years from 1990 to 2018 of gas stations holding at least one tank active at the December 23, 1998 reform, in Texas and the 23 control states. Vertical bars are 95 percent confidence intervals from standard errors clustered by state.

facility discovers 0.55 fewer releases at closure per 100 facility-years (s.e. 0.08), against a control-state rate of 0.58. Texas closed more tanks after the reform, and the tanks it closed were consistently cleaner when inspected.

Panel (b) shows releases discovered outside a closure. The coefficients are close to zero in the first years after the reform and rise to between 0.5 and 0.8 points from 2011 to 2016, against a control-state rate of 2.1 percent. The individual estimates are imprecise, and this late increase is suggestive rather than settled.

5 Dynamic Structural Model

I build a dynamic model of how gas stations manage their underground tanks so that I can compare the instruments states use to retire old tanks. States act on the same aging capital in four ways. Texas prices coverage by the risk of each tank, and the fund states charge one flat fee per tank. Some states pay for replacement directly. California’s RUST program gives small operators grants and below-market loans to replace single-walled tanks with double-walled ones ([California State Water Resources Control Board 2025](#); [State of California 2026](#)). Other states set a deadline. California required every single-walled tank to be permanently closed by the end of 2025 ([State of California 2014](#)), Kentucky required bare-steel tanks to close from 2012 ([Kentucky Energy and Environment Cabinet 2011](#)), and Maine phased out unprotected steel tanks between 1989 and 1997 ([State of Maine 2026](#)). [Section 4](#) measured the effect of one of these changes, the move in Texas from a flat fee to risk-based prices. That estimate cannot rank the other instruments. A subsidy acts on the cost of installing a new tank, a mandate removes the option to keep an old one, and no reform in my data moves either margin. The model recovers the parameters that stay fixed when the policy changes: the value a facility places on operating revenue, the weight it places on a dollar of coverage cost, and the costs of removing tanks, installing them, closing, and opening a new station. With those held fixed, [Section 6](#) puts Texas on a flat fee and the fund states on risk-rated premiums, adds a replacement subsidy and a removal mandate, and compares each with a first-best premium. A below-market loan lowers the effective cost of a new tank, which is the margin the subsidy experiment moves.

The model is built around one decision. A station owner holds tanks that grow older every year, and every year the owner pays to keep them: the insurance premium, plus the deductible owed if a tank leaks times the chance that one does. That recurring charge rises as the tanks age. The owner can instead pay a large bill once, to remove the oldest tanks, to replace them with new double-walled tanks, or to close the station and sell what remains. Paying now lowers the charge in every year that follows. Waiting keeps the money and leaves the owner free to decide again next year. Which choice wins depends on how much the station earns and how many more years the owner expects to run it. A station with decades of revenue ahead spreads the one-time bill over many years of lower charges, and a station near the end of its life cannot. The policies differ in which part of this decision they touch. A risk-based premium raises the recurring charge on old tanks, a subsidy lowers the one-time bill for a new tank, and a mandate takes away the choice to wait. Someone deciding whether to open a new station faces the same terms from the other side: open if running a station with new tanks under the state’s rules is worth what it costs to build one.

5.1 The model

The model has two kinds of agents. Incumbent facilities decide each year what to do with the tanks they hold, and potential entrants decide whether to open a new station. Both face the same insurance rules. I develop the incumbent’s problem first, because an entrant that opens becomes an incumbent and the value of opening is the value of that problem.

5.1.1 The incumbent and its state variables

The incumbent is the facility rather than the tank. Closure in the data is a portfolio decision, several tanks removed in one act, and a model of one representative tank cannot price the partial-closure margin [Section 4.3](#) found. A facility holds a fleet of tanks that differ in age and wall construction. It earns revenue on the fuel its tanks store, pays the premium its insurer sets on the fleet, and bears the deductible if a leak arrives. As the tanks age, the chance of a leak rises with them, and the expected out-of-pocket cost therefore climbs under any pricing rule. The premium is the part the pricing rule decides. Under a risk-rated schedule it climbs with the tanks, and under a flat fee it does not move with age at all.

An incumbent's state variables are $s = (n, B, g)$. The state g is the U.S. state in which the facility operates. It sets the fuel margin the facility earns and the insurance rules it faces: the premium schedule P_g and the deductible D_g . In a fund state, P_g is the fund's fee per tank. In Texas it is the filed rate schedule of the carrier writing the facility's coverage that year. Facilities in the data switch carriers, and the model takes each year's carrier as given rather than pricing the choice of carrier. The capacity class B records how much fuel the facility can store. The fleet n records the tanks themselves.

The fleet n describes the facility's tanks: how many it holds, how old each one is, and whether each is single- or double-walled. The premium and the hazard depend on the whole fleet, $P_g(n)$ and $H(n)$, and not only on how many tanks it holds. I write $N(n)$ for the tank count. Capacity is kept apart from the fleet because the two set different payoffs. Capacity sets revenue, and the count, age, and wall construction of the tanks set the premium and the hazard.

The composition of the fleet matters because the premium and the hazard are steep in exactly the tanks a removal can take out. Consider two facilities in the same state under the same filed schedule. The first holds three tanks, two old single-walled tanks and one new double-walled tank. The schedule charges the old single-walled pair the highest per-tank rate on the card and the new tank the lowest, and the facility's hazard sits high on the risk gradient of [Section 3.1](#). Removing those two tanks is one action that eliminates most of the facility's premium and moves its hazard to the bottom of the gradient. The second facility holds four double-walled tanks of middle age. Its total premium is larger because it insures more tanks, but every tank carries the same low rate, its hazard sits near the bottom of the gradient, and no single removal moves either charge by much. A tank count calls the second facility the larger risk, and an average age calls the two similar. Tracking the fleet is what lets the model see that the first facility can buy a large reduction in both charges with one removal and the second cannot.

Within a year, events run in a fixed order. The facility observes its state variables, draws one private payoff shock for each action available to it, and acts. The action reshapes the fleet at once, the insurer prices the fleet that remains, the year's payoff is realized on that fleet, and the surviving tanks then age into next year's state variables.

5.1.2 Choices and how the fleet moves

Each year the facility chooses $d = (k, m)$, to remove k tanks and install m , or it chooses to exit, $d = X$. Maintaining is $(0, 0)$. Downsizing removes tanks without installing any, and replacement removes and installs in the same year. Exit is permanent. Every tank installed is new and double-walled, mirroring the secondary-containment requirements that ended single-walled installation over the sample. The action leaves a fleet n' in the ground. Every surviving tank then ages whether or not the facility acted. The facility cannot slow its tanks' aging. Its only levers on the fleet's risk are removal and installation, and those are the margins the premium prices. I write $F_d(s' | s)$ for the probability of reaching s' after action d at s : the reshaping of the fleet by the action, followed by aging.

5.1.3 Payoffs

An operating action pays the facility its revenue less what it costs to carry its risk and to change its fleet,

$$u_{(k,m)}(s) = \underbrace{\pi(B, g)}_{\text{operating revenue}} - \underbrace{C_g(n')}_{\text{coverage cost}} - \underbrace{A(k, m)}_{\text{adjustment cost}}, \quad u_X(s) = \underbrace{X(n)}_{\text{scrap value}}.$$

Operating revenue $\pi(B, g)$ is what a facility of capacity B earns selling fuel in state g . The coverage cost is what the facility pays each year for carrying its risk,

$$C_g(n') = P_g(n') + H(n') D_g,$$

the premium the insurer charges on the post-action fleet plus the expected out-of-pocket cost: the chance $H(n')$ that the facility has a leak this year times the deductible D_g . Both parts are dollars the facility pays, and the pricing rule sets the first. The adjustment cost $A(k, m)$ is what it costs to remove k tanks and install m . Exit pays the scrap value $X(n)$ of the fleet and nothing afterward.

Each year the facility also draws a private shock ε_d for every action, and its total payoff from action d is $u_d(s) + \varepsilon_d$. The shocks carry the value that no measured input records, a lease about to end or a site the owner values for reasons of their own, and they are why two facilities with identical state variables choose differently. All else equal, a facility whose coverage cost is concentrated in a few old single-walled tanks gains more from removal than one whose cost is spread evenly across its fleet, and a steeper premium schedule widens that difference.

5.1.4 The incumbent's decision rule

The facility knows its premium schedule, its hazard, its deductible, and this year's shocks. It treats the rules it faces as permanent and forecasts the aging and leaking of its tanks with the same processes the model estimates. It does not observe future shocks. The facility takes the action with the highest value,

$$d^*(s, \varepsilon) = \arg \max_{d \in \mathcal{D}(s)} \{v_d(s) + \varepsilon_d\}, \quad v_d(s) = u_d(s) + \beta \sum_{s'} F_d(s' | s) V(s'),$$

where $\mathcal{D}(s)$ is the set of actions available at s , β is the discount factor, and $V(s)$ is the expected value of starting a year at s before the shocks are drawn. Exit carries no continuation value. The shocks are independent Type-I extreme value, and the value function and choice probabilities therefore take the logit form,

$$V(s) = \log \sum_d \exp v_d(s), \quad \Pr(d | s) = \frac{\exp v_d(s)}{\sum_{d'} \exp v_{d'}(s)}.$$

The dynamic tradeoff is the core economics of the model. Keeping an old tank means paying its premium and bearing its hazard again next year. Removing it means paying the adjustment cost once. The choice turns on how long the facility expects to keep operating. The same repricing therefore moves a facility near the end of its horizon differently than one with decades of revenue ahead, and a static model of the same choice cannot represent that difference.

5.1.5 Entry

New stations open each year, and the same insurance rules shape that decision. In each state there are people or firms who could open a station in a given year. Opening one means buying a set of new double-walled tanks and paying a one-time cost of opening: the site, the permits, the building, and the dispensers. The return is the value of running the station under the state's insurance rules for as long as it pays to do so, which is the value of an incumbent that starts with a new fleet. A potential entrant opens when that value covers the cost. The rules matter because they set what a new fleet pays every year to carry its risk.

Each year a pool of potential entrants considers opening in state g . A potential entrant opens if

$$\underbrace{\bar{V}_g}_{\text{value of a new facility}} - \underbrace{A(0, \bar{n}_g)}_{\text{cost of its tanks}} - \underbrace{ec}_{\text{cost of opening}} + \varepsilon^E \geq \varepsilon^O,$$

and stays out otherwise. $\bar{V}_g$ is the expected value of an incumbent in state g that starts the year with a new double-walled fleet, averaged over the sizes at which entrants open there. $\bar{n}_g$ is the number of tanks an entrant installs, and $A(0, \bar{n}_g)$ is the cost of installing them, at the same install cost an incumbent pays. The entry cost ec is everything else it costs to open a station, paid once. The shocks ε^E and ε^O are the entrant's private values of opening and of staying out, independent Type-I extreme value with the same unit scale as the incumbent's shocks. The probability that a potential entrant opens is therefore

$$p_g = \Lambda(\bar{V}_g - A(0, \bar{n}_g) - ec), \quad \Lambda(x) = \frac{1}{1 + e^{-x}}.$$

Once open, an entrant is an incumbent and solves the problem above. Each potential entrant compares only its own value with its own cost, and the number of others who open does not enter its decision. A state's insurance rules reach the entry decision only through $\bar{V}_g$. A premium schedule that charges a new double-walled fleet little raises the value of opening, and a fee that charges it the same as an old fleet does not. The cost of opening, $A(0, \bar{n}_g) + ec$,

is the same under every policy. A replacement subsidy pays for tanks that incumbents install and does not lower the cost of a new station’s tanks.

5.1.6 Where policy enters the model

The four instruments states use each change one object in this model and leave the rest fixed. Risk-based pricing, as Texas has used since its reform, and a flat fee per tank, as the fund states charge, differ only in the premium schedule $P_g(\cdot)$. A flat fee makes P_g depend on the tank count $N(n)$ alone, and the part of the coverage cost a facility can lower by choosing which tanks to keep shrinks to the deductible term. A risk-rated schedule makes P_g steep in age and wall construction, and removing the right tank then lowers both the premium and the expected out-of-pocket cost. A replacement grant, such as California’s RUST grants, or a below-market replacement loan lowers the install part of $A(k, m)$ for incumbents and leaves the entrant’s cost of opening unchanged. A removal deadline, such as California’s requirement that single-walled tanks close by the end of 2025, changes no payoff. It removes from $\mathcal{D}(s)$ every action that leaves a tank past the deadline in the ground. Under each rule incumbents re-solve their problem, and new stations open at the rate the new $\bar{V}_g$ implies. The counterfactuals of [Section 6](#) are these substitutions performed on the estimated model.

5.2 Estimation and identification

To estimate the model, I first build every price, risk, and revenue input from data outside the dynamic model and hold those inputs fixed. I then estimate the payoff parameters by nested pseudo-likelihood ([Aguirregabiria and Mira 2007](#)), which chooses the parameters that maximize the likelihood of the actions facilities take, and I estimate the entry cost from the number of stations that open each year given the solved incumbent problem. [Section D.2](#) gives the solver and the convergence criteria. I first describe how the model is parameterized and then what variation in the data identifies each parameter.

The estimating sample is the incumbent panel at annual frequency, 1,654,626 facility-years. Texas enters from 2006 through 2020, with the start set by when its carrier premium data begin, and the control states enter from 1999 through 2020. Control states number 23 in the reduced form, and 14 enter estimation. Kansas and Maryland are excluded because no per-tank fund fee could be verified for either, and an unverifiable fee is missing data rather than a price. Several control states charge no per-tank fee at all, and those verified zeros enter the model as true prices. Missouri is excluded for holding too few events to identify a state level. Standard errors come from the outer product of the score at the estimates.

5.2.1 Parameterization

I group tank ages into eight five-year bands, 0 to 4 years, 5 to 9, and so on through 30 to 34, with an eighth band for tanks 35 years and older. The fleet is then the 16 counts $n_{w,b}$ of tanks of wall type w , single-walled or double-walled, in age band b , and the tank count is $N(n) = \sum_{w,b} n_{w,b}$. A facility holds between one and six tanks. Facility-years that begin with more than six tanks are left out of estimation. Enumerating every fleet of one to six tanks gives 74,612 fleets, and crossing them with the four capacity classes gives 298,448

combinations of fleet and capacity within each state. The value function solves exactly at these sizes. With logit shocks, the value function at any parameter is the solution of one linear system, and no approximation is needed.

The model limits what a facility can do in one year in four ways. First, a facility removes at most four tanks in a year. Second, a facility that stays open keeps at least one of its existing tanks through the removal. Removing every tank is exit, and a facility with N tanks that does not exit can therefore remove at most $N - 1$ of them. Third, a facility installs at most four tanks in a year, and the fleet after removal and installation can hold no more than six. Fourth, a facility installs only in a year in which it also removes. Exit is available in every state. Observed actions are coded onto this menu. A year in which a facility removes more than four tanks is coded as removing four. A year in which a facility removes every tank and installs new ones is coded as removing all but one. Years in which a facility only adds tanks, and years whose action would leave more than six tanks, are left out of estimation.

A rule decides which tanks a removal takes. It removes single-walled tanks first, oldest band first, and then double-walled tanks, oldest band first. Installed tanks enter the youngest double-walled band. Each surviving tank below the oldest band then advances one band with a probability a_b that depends only on its current band, and otherwise stays. I estimate a_b as the share of tank-years in the registry whose age crosses into the next band within the year (Table 15). The oldest band is absorbing. The capacity class moves too, because removing and installing tanks changes how much fuel a facility can store. Its transition probabilities are the observed shares of consecutive facility-years that move from one class to another. I estimate them separately by the net change in tanks the action made, $m - k$, grouped as three or more fewer tanks, two fewer, one fewer, no net change, and one or more additional tanks, and for years with no action. The removal rule, the aging probabilities, and the capacity transitions together give $F_d(s' | s)$.

Revenue in year t is

$$R_{B,g,t} = \frac{\bar{c}_B \cdot \bar{m}_{g,t}}{\$10,000},$$

where $\bar{c}_B$ is the mean storage capacity of capacity class B , in gallons, and $\bar{m}_{g,t}$ is the fuel margin in state g : the retail gasoline price less the wholesale price, the state gasoline tax, and the federal tax. I average the margin over the rate era containing year t , 1999 to 2013, 2014 to 2018, or 2019 to 2020. A facility therefore forecasts a steady local margin rather than a yearly swing. Each tank's capacity is capped at 60,000 gallons before a facility's capacity is summed, and the four classes are the quartiles of that capped capacity. Revenue is measured outside the model, following Sweeting (2013), and enters the payoff through one estimated weight. Section D.1 gives the price sources.

The premium a Texas facility pays in year t comes from the filed rate schedule of its carrier c that year,

$$P_{g,t}(n) = \max \left\{ \delta_{c,t}(N(n)) \sum_{w,b} n_{w,b} p_{c,t}(w,b), \underline{P}_{c,t} \right\},$$

where $p_{c,t}(w,b)$ is the premium the carrier's rate card charges one tank of wall type w in age band b , $\delta_{c,t}(N)$ is the carrier's discount for insuring N tanks at one facility, and $\underline{P}_{c,t}$

is its minimum premium. Every Texas premium is priced at the same reference contract, \$1,000,000 in coverage with a \$5,000 deductible (Section A.6). In a fund state the premium is the fund's fee per tank times the tank count, $P_g(n) = f_g N(n)$. The hazard $H(n)$ is the annual probability that the facility reports a leak. I read it from the leak model of Section 3.1 at the fleet's majority wall type and at the age band of its average tank age, weighting each tank equally. A fleet with as many single-walled tanks as double-walled tanks is rated as single-walled. The deductible D_g is the fund's statutory deductible in a fund state and \$5,000 in Texas.

The payoff is linear in the parameters,

$$u_{(k,m)}(s) = \psi R_{B,g,t} + \alpha_g - \gamma [P_{g,t}(n') + H(n') D_g] - k c^{\text{rem}} - m c^{\text{inst}}, \quad u_X(s) = \kappa_1 N(n).$$

The weight ψ converts measured revenue into payoff. The state effect α_g absorbs what is persistently different about operating in state g , with Texas the reference, $\alpha_{\text{TX}} = 0$. The weight γ applies to one dollar of coverage cost, premium and expected out-of-pocket cost alike. The costs c^{rem} and c^{inst} are paid per tank removed and per tank installed, and the adjustment cost is therefore $A(k, m) = k c^{\text{rem}} + m c^{\text{inst}}$. Exit pays a scrap value of κ_1 per tank, with the constant part of exit value normalized to zero. The shock scale is normalized to one, and every parameter therefore reads in units of the shock. The discount factor is calibrated at $\beta = 0.95$. One parameter vector applies in every state and every Texas carrier-year.

5.2.2 Entry parameterization

Entry needs a number of potential entrants in each state and year, and that number is not observed. I set it by rule. Let $M_{g,t}$ be the number of facilities operating in state g at the start of year t , and let $E_{g,t}$ be the number that open during the year. I take the most facilities per resident observed in any state and year of the entry sample, $\lambda = \max_{g,t} (M_{g,t} + E_{g,t}) / \text{POP}_{g,t}$, as the most a state's population can support. The search leaves out Maine and South Dakota, whose registry counts are not reliable enough to define a maximum, and Section D.6 names the state and year that set it. The number of potential entrants is that cap less the facilities already operating, $\text{Pool}_{g,t} = \lambda \text{POP}_{g,t} - M_{g,t}$.

An entrant's size is taken from the data. Entrants open at the tank counts and capacity classes observed among entrants in their state, every tank new and double-walled. That gives the share $\omega_g(s)$ of entrants that start at each fleet and capacity class and the mean tank count $\bar{n}_g$. The value of a new facility is the incumbent's value at those starting points, averaged with those shares, $\bar{V}_g = \sum_s \omega_g(s) V_g(s)$, evaluated at the estimated incumbent parameters. The probability that a potential entrant opens is $p_g = \Lambda(\bar{V}_g - c^{\text{inst}} \bar{n}_g - ec)$, the decision rule of Section 5.1 with the install cost per tank. The term $c^{\text{inst}} \bar{n}_g$ does not change under any counterfactual. The number that open in a state and year is then binomial with $\text{Pool}_{g,t}$ trials and probability p_g , and I choose ec to maximize that likelihood.

The entry sample is 259 state-years in 14 states, the control states from 1999 through 2017 and Texas from 2006 through 2017, in which 12,888 facilities open. The window ends in 2017 because the population series does. A state-year whose pool is smaller than its own count of entrants is excluded. South Dakota's recorded stock exceeds its cap in every year, and it is

left out of the fit. The standard error on ec is clustered by state. It treats $\bar{V}_g$ as known. The value of a new facility is built from the estimated incumbent parameters, and the standard error does not carry their sampling error. It therefore understates the uncertainty in ec .

5.2.3 Revenue weight

The revenue weight ψ is identified from how exit and downsizing rates differ across facilities whose measured revenue differs. Measured revenue varies for two reasons. Larger capacity classes store more fuel, and the fuel margin differs across states and eras. Both components come from public data: the margin from EIA retail and wholesale prices and published tax rates, and capacity from the tank registry. Neither draws on a facility’s own sales or accounts. The margin is a state-level market price that no single station sets, and the capacity class that sets this year’s revenue was fixed before this year’s shocks were drawn. A removal or installation this year changes revenue only from next year onward. Facilities with higher measured revenue exit and downsize less often, and ψ is the weight that reproduces how strongly their choices move with it. [Section 5.3](#) discusses what the constructed measure means for how ψ is read.

5.2.4 Coverage-cost weight

The coverage-cost weight γ is identified from variation in what a facility pays to carry its risk that the risk itself does not explain. I compute that cost for every facility, every year, and every action it could take. The Texas premium comes from running the facility’s post-action fleet through the reimplemented rate engine of its carrier in that year, at the reference contract. The fund premium is the fund’s fee per tank times the tank count. The expected out-of-pocket cost is the hazard from the leak model times the state’s deductible. Three comparisons separate the price of risk from the risk itself. First, the Texas premium rises with the age of the fleet, and the fund fee does not move with age at all. Second, within Texas, carriers file different schedules, and identical fleets in the same year pay different premiums when different carriers write them. Third, fund deductibles differ from state to state at a given premium, and each fund sets its deductible by rule for every facility it covers. Carriers and funds set these schedules for every facility in a rating class rather than for an individual facility, and a facility’s private shock cannot move them. Facilities carrying an imputed market-mean premium contribute nothing to the second comparison. [Figure 34](#) and [Figure 35](#), in [Section D.4](#), plot the first and third comparisons, the premium and the expected out-of-pocket cost by fleet age in each state.

5.2.5 Adjustment costs

The removal and install costs are identified from how often facilities act given what acting saves them. For every fleet, the first-step inputs give the coverage cost before and after each removal and each replacement, and that saving differs across states, carriers, years, and fleets for reasons set by the schedules and the hazard rather than by the facility. Facilities whose fleets would save more by removing a tank downsize more often. Given γ , which converts the saving into payoff, c^{rem} is the per-tank cost that reproduces how often facilities remove tanks at each level of saving. The install cost is read from the choice between the two ways of

removing a tank. A facility that removes old tanks can stop there or put new double-walled tanks in their place, which keeps its tank count and capacity and carries the lowest coverage cost on the schedule. The share of removals that come with installations pins c^{inst} . Timing driven by forces outside the model, the end of a lease for one, adds noise to these estimates rather than bias, provided it is uncorrelated with the saving.

5.2.6 State effects and scrap value

The state effects α_g are identified from the operate-versus-exit margin. A state's α_g is read off how often its facilities keep operating, relative to Texas facilities holding the same fleet at the same coverage cost. The scrap value κ_1 is identified from how the exit rate changes with tank count within a state. Both touch exit, and they separate by shape: a level per state against a slope in tank count. No state's level reaches a bound in the estimation.

5.2.7 Entry cost

The entry cost is identified from the number of stations that open relative to the pool. Given the pool and the value of a new facility in each state, ec is the cost at which the predicted number of entrants matches the observed number across all state-years. One parameter reproduces the total number of entrants and does not reproduce their distribution across states ([Section D.6](#)). Its level depends on the size of the pool, which [Section 5.3](#) discusses.

5.3 Discussion of model assumptions

5.3.1 Single agent

No facility in the model responds to another. The reform repriced each facility's own risk rather than the structure of local competition, and the responses the model must reproduce are choices over a facility's own capital: removing, replacing, and exiting. The incumbent's problem is therefore a single-agent dynamic discrete-choice problem in the tradition of Rust ([1987](#)). The cost of that scope is that the model is silent on how pump prices and fuel volumes adjust as stations exit and enter. The return is that the state variables reduce to one facility's own fleet, capacity, and state, and the problem solves without approximation.

5.3.2 Revenue measured outside the model

Revenue enters as a measured lookup, and a facility changes its revenue only by changing its capacity class, through removal and installation choices the model already prices. I observe no facility's own gallons sold and no facility's own prices, and the measure is built from what is observable instead, the capacity class and the state fuel margin. Two consequences follow. The measure does not observe how many times a facility turns over its storage in a year, and ψ absorbs that turnover rate. And the measure assigns every facility in a class, state, and era the same revenue, and true differences within a cell enter the estimate as measurement error, which attenuates ψ toward zero. Both forces push the same way. The estimate should therefore be read as a lower bound on how revenue moves choices, a weight

on the constructed measure rather than the value of a literal dollar. The idiosyncratic part of profitability is carried by the payoff shocks rather than measured.

5.3.3 The carrier is taken as given

The carrier writing a Texas facility's coverage is taken as given. If facilities switch toward carriers that price their fleet more cheaply, the premium a facility pays is partly chosen, and γ understates the response to a premium a facility could not shop away.

5.3.4 The Texas deductible

The financial-responsibility contract records carry neither the deductible nor the coverage form a facility chose, and the rate engines therefore price every Texas contract at the carriers' filing base, \$1,000,000 in coverage with a \$5,000 deductible. I impose that deductible on every Texas facility. The assumption makes the out-of-pocket cost move within Texas only through the hazard. Its cost is measurement error in the out-of-pocket level, because actual deductibles that vary across Texas facilities differ from the imposed one, and if the choice of deductible tracks a facility's risk the error is not random. For that reason the cross-state variation in fund deductibles carries the identification of the level.

5.3.5 Which tanks are removed

The facility chooses how many tanks to remove, and the rule that removes the oldest single-walled tanks first chooses which. The data show the margin is nearly deterministic. The rule names the removed category correctly in 84.6 percent of removal events ([Appendix D](#)). Where the rule misnames the removed tank, the premium and the hazard carry small assignment errors.

5.3.6 Linear adjustment costs and the action limits

Removal and installation costs are linear in the number of tanks, and acting carries no fixed cost. A facility that removes three tanks pays three times what one removal costs, and the model must explain years in which several tanks move at once by the coverage cost each removal saves. If acting carries a fixed cost, such as a site assessment paid once whatever the number of tanks, the per-tank costs absorb it. The action limits also misstate a small set of years. Removals of more than four tanks are coded as four, and a year in which a facility replaces every tank is coded as keeping one old tank.

5.3.7 Stationary rules and a shared forecast

A facility treats the rules it currently faces as permanent. A rule change that a facility in fact foresaw would be read by the model as a response to current prices. Facilities also forecast their tanks' aging and leaking with the same processes the first step estimates, and the facility and the model therefore share one forecast.

5.3.8 The discount factor and the normalizations

The discount factor is calibrated at $\beta = 0.95$, because choice data cannot separate patience from the flow payoffs. The shock scale is normalized to one, and every parameter reads in shock units, with ratios of parameters carrying the economic content. Fixing the shock scale and estimating ψ is the same normalization Sweeting (2013) makes from the opposite side, holding measured revenue at par and estimating the shock scale instead. The constant part of exit value is normalized to zero, which pins the level of the value function. The data therefore identify how exit value scales with tank count and not its level.

5.3.9 Entry

Entry rests on three assumptions. Entrants do not respond to one another, and each compares its own value of opening with its own cost. The pool of potential entrants is set from the highest observed number of facilities per resident rather than estimated, and the level of ec depends on it. A larger pool implies a lower entry probability for the same count of entrants and therefore a higher entry cost. Section D.6 reports the estimate with pools two and five times as large and with no limit on the pool. Finally, facilities open at the sizes observed among entrants, and the model does not price the choice of size.

5.3.10 Cross-sectional identification

Identification is cross-sectional throughout. The parameters are pinned by dispersion across facilities and states within years, not by the before-and-after of the Texas reform, whose pre-period never enters the estimation window. The estimates hold conditional on the assumptions above, and Section 6 inherits them.

6 Estimates, Counterfactuals, and Welfare

In this section, I use the estimated model to evaluate how the pricing of tank insurance affects tank closure, leaks, and welfare. I start by presenting the parameter estimates and the model's fit to observed choices. I then run four counterfactuals on the observed stock of facilities. The first two swap the two pricing regimes in the data, putting Texas on a flat fee per tank like the fund states and putting the fund states on risk-rated premiums like Texas. The other two add the replacement subsidy and the removal mandate that U.S. states use to retire old tanks. I rank each against a first-best premium that charges each facility the full social cost of its leak risk.

6.1 Estimates and model fit

Table 2 reports the five payoff parameters and the entry cost, with standard errors. Dollar figures are payoff units divided by γ and multiplied by \$10,000, and they read as premium-equivalent dollars. Facilities respond to the coverage cost. The weight on a dollar of premium or expected out-of-pocket cost is $\gamma = 0.550$ (s.e. 0.007). The weight on measured revenue is $\psi = 0.072$ (s.e. 0.002). It is low as the construction of the revenue measure anticipates,

because ψ absorbs the turnover normalization and the attenuation [Section 5](#) describes, and I read it as a lower bound on how revenue moves choices.

The adjustment costs are large. Removing a tank costs $c^{\text{rem}} = 4.906$ (s.e. 0.009), about \$89,142, and installing one costs $c^{\text{inst}} = 0.931$ (s.e. 0.011), about \$16,911. These are full project costs rather than the price of a tank, and their size is what rationalizes how rarely facilities act. On the total they land where engineering estimates do. Removing a tank and installing its replacement costs about \$106,054 in the model. One Indiana county’s project under the state trust fund put removing and replacing three tanks at \$133,000 to \$167,000 per tank ([Indiana Capital Chronicle 2026](#)), the EPA’s 1988 estimate implies \$65,000 to \$86,000 per installed tank in 2023 dollars ([EPA 1988](#)), and a 1987 industry survey puts an installed double-walled tank at \$98,000 to \$105,000 ([U.S. General Accounting Office 1992](#)). Clean removal costs \$10,000 to \$25,000 per tank in state program records and reimbursement schedules ([Association of State and Territorial Solid Waste Management Officials 2017](#); [Kentucky Energy and Environment Cabinet 2024](#)), while the model puts removal at about 5 times the install. The benchmarks price the tank work, and the model’s costs price the whole project. Every action that opens the ground pays c^{rem} , and the parameter therefore carries the common cost of acting: labor, inspections, environmental review, downtime, financing, and above all the closure-triggered site assessment. Contractor guidance for an old single-walled tank is to budget two to three times the clean-closure cost against contamination found on removal ([UST Contractors 2026](#)). The claims records put the mean cleanup at \$348,000 once a release is found, and an expected discovery cost of about 21 percent of that cleanup closes the gap between the middle of the clean-removal range and the estimate. Read against that, c^{inst} is the incremental cost of setting a new tank in a hole that is already open. The 1987 and 1988 figures and the removal range are in 2023 dollars, and the Indiana figure is in the dollars of its own year. The benchmarks are a plausibility check on magnitudes rather than an independent estimate of either parameter.

Leaving is worth little, and opening is expensive. The exit value is $\kappa_1 = 0.047$ (s.e. 0.010), about \$859 of scrap per tank, and what disciplines exit is therefore forgone operating value rather than salvage. The entry cost is $ec = 12.71$ (s.e. 0.89), about \$230,944 to open a station before the install cost of its tanks. The state effects are estimated alongside and not reported. The iteration converges in the parameters at iteration 13, and [Section D.2](#) documents the convergence diagnostics.

The model reproduces the observed action shares, which the estimation touches only through the likelihood and does not match by construction. Across the estimation sample it predicts maintaining in 97.4 percent of facility-years against 97.8 percent observed, and exit in 1.6 percent against 1.6 percent. In the control states every predicted share lies within half a percentage point of the observed one. Texas is where one parameter vector strains. There the model predicts downsizing in 1.1 percent of facility-years against 0.1 percent observed and exit in 0.1 percent against 0.7 percent. It attributes part of the Texas response to the partial margin the premium prices most directly rather than to full exit.

Table 2: Structural parameter estimates

Parameter	Description	Estimate	Standard error	Dollar equivalent
ψ	Weight on measured revenue	0.072	(0.002)	
γ	Weight on coverage cost	0.550	(0.007)	
c^{rem}	Removal cost, per tank	4.906	(0.009)	\$89,142
c^{inst}	Install cost, per tank	0.931	(0.011)	\$16,911
κ_1	Exit value, per tank	0.047	(0.009)	\$859
ec	Entry cost, per facility	12.711	(0.893)	\$230,944

Notes: Estimates of the payoff parameters of the dynamic model of [Section 5](#) and of the entry cost, with standard errors in parentheses. One model unit is 10,000 dollars per facility-year. The dollar equivalent of a cost or a value is its estimate divided by γ , the weight on a dollar of coverage cost, and multiplied by 10,000. The weights ψ and γ have no dollar equivalent. The removal and install costs are full project costs per tank, covering labor, inspections, environmental review, and replacement capital. The five payoff parameters are estimated by nested pseudo-likelihood on 1,654,626 facility-years, Texas from 2006 to 2020 and 14 control states from 1999 to 2020, with standard errors from the outer product of the score. The control-state effects are estimated and not reported, and Texas is the omitted state. The entry cost is estimated in a separate binomial step on 259 state-years, at the baseline pool of potential entrants, with a standard error clustered by state that takes the value of a new facility as given. The entry cost excludes the install cost of the entrant’s tanks.

6.2 Counterfactuals

Each counterfactual changes one object of the model and leaves the rest fixed. The first two swap the two pricing regimes in the data. (CF1) puts Texas on a flat fee per tank, the way the fund states price coverage, and leaves the fund states unchanged. (CF2) puts the fund states on risk-rated premiums, the way Texas prices coverage, and leaves Texas unchanged. The other two impose policies U.S. states use to retire old tanks, on top of the observed contracts. (CF3) is a grant toward the cost of new tanks, as in Arizona’s Tank Site Improvement Program ([State of Arizona 2019](#)), California’s replacement grants ([California State Water Resources Control Board 2025](#)), and the reimbursements Kansas and Louisiana pay for replacing single-walled systems with double-walled ones ([Kansas Department of Health and Environment 2026](#); [Louisiana Department of Environmental Quality 2008](#)). (CF4) is a removal deadline. Massachusetts required single-walled steel tanks to close by August 2017 ([Massachusetts Department of Environmental Protection 2017](#)), California requires single-walled tanks to close by the end of 2025 ([State of California 2014](#)), and Maine retired unprotected steel tanks on an age schedule that ended in 1997 ([State of Maine 2026](#)). A benchmark, (FB), is a premium equal to the full social cost of each facility’s leak risk, which no state charges. Each counterfactual holds the estimated parameters fixed, changes one policy, and re-solves the incumbent’s problem and the entry rule. The state effects α_g stay in both the facility’s re-optimization and the welfare level, because they stand for persis-

tent features of operating in a state rather than the policy under study. Welfare starts from the observed stock of facilities. The baseline is the observed contracts, risk-rated pricing in Texas and a flat fee in every fund state, and it reproduces observed behavior. Under it a maintaining facility earns

$$u_{(0,0)}(s) = \psi R_{B,g,t} + \alpha_g - \gamma [P_{g,t}(n') + H(n') D_g], \quad P_{g,t}(n') = f_g N(n') \text{ for } g \neq \text{TX}. \quad (\text{CF0})$$

(CF1) Flat fee in Texas. The first counterfactual replaces the filed Texas premium schedules with a flat fee of \$50 per tank, the annual fee per tank that Texas charged tank owners under its state fund program. The Texas deductible stays at \$5,000, and the fund states keep their own fees and deductibles. A maintaining Texas facility then earns

$$u_{(0,0)}(s) = \psi R_{B,\text{TX},t} - \gamma \left[\underbrace{\tau N(n')}_{\tau = \$50 \text{ per tank}} + H(n') D_{\text{TX}} \right], \quad \alpha_{\text{TX}} = 0. \quad (\text{CF1})$$

The Texas fund also collected a fee on fuel deliveries, which the data do not record, and the counterfactual leaves it out. Revenue, the hazard, and every cost hold fixed, and only the price of risk changes. Risk-based pricing means Texas in this sample, and the contrast identifies the pricing mechanism suggestively rather than cleanly.

(CF2) Risk-rated pricing in the fund states. The second counterfactual is the reverse. It keeps the filed Texas schedules in Texas and gives every fund state the Texas premium for each of its tanks. A tank pays $\bar{p}^{\text{TX}}(w, b)$, the premium Texas carriers charge one tank of wall type w in age band b , averaged over carriers and years with Texas tank-years as weights, and the deductible is the Texas \$5,000. A maintaining facility in a fund state then earns

$$u_{(0,0)}(s) = \psi R_{B,g,t} + \alpha_g - \gamma \left[\underbrace{\sum_{w,b} n'_{w,b} \bar{p}^{\text{TX}}(w, b)}_{\text{Texas cell premium replaces } f_g N(n')} + H(n') \underbrace{D_{\text{TX}}}_{\$5,000} \right], \quad g \neq \text{TX}. \quad (\text{CF2})$$

The multi-tank discount and the minimum premium of the Texas schedules are dropped, and a facility pays the sum of its tanks' cell premiums. Revenue, the hazard, and every cost hold fixed.

(CF3) Replacement subsidy. The third counterfactual pays a share $\rho = 50$ percent of the install cost of every tank an incumbent installs, in every state and under the observed contracts. An incumbent that removes k tanks and installs m earns

$$u_{(k,m)}(s) = \psi R_{B,g,t} + \alpha_g - \gamma [P_{g,t}(n') + H(n') D_g] - k c^{\text{rem}} - m \underbrace{(1 - \rho) c^{\text{inst}}}_{\text{subsidized install}}, \quad (\text{CF3})$$

and the government pays the remaining ρc^{inst} on each tank. That payment is the one direct outlay in the accounting below. The subsidy does not lower the cost of a new station's tanks, and it reaches entry only through the value of operating, $\bar{V}_g$.

(CF4) Removal mandate. The fourth counterfactual requires that a single-walled tank past twenty years be removed, in every state and under the observed contracts. Twenty is where the single-walled release hazard peaks ([Section 3.1](#)), and filed premiums imply that

replacement beats deferral for single-walled systems near that age (Yin, Kunreuther, and White 2011). Because a mandate at any age retires tanks on age alone, the cost reported below comes from what a fixed age cutoff does rather than from where it is set. The mandate changes no payoff. It removes from the choice set every action that leaves such a tank in the ground,

$$\mathcal{D}^{\text{mandate}}(s) = \{ d \in \mathcal{D}(s) : n'(d) \text{ holds no single-walled tank past age twenty} \}, \quad (\text{CF4})$$

and a facility meets it by downsizing, replacing, or exiting at the estimated costs.

(FB) First-best premium. The benchmark charges every facility, in every state, a risk-rated premium equal to the marginal social cost of its own leak risk,

$$P_g^{\text{fb}}(n') = H(n') [S(n') + \bar{S}_{>L} + E_{\text{ext}}], \quad (\text{FB})$$

the post-action hazard times the sum of the expected claim the insurer pays, $S(n')$, the expected cleanup above the coverage limit, $\bar{S}_{>L}$, and the damage a leak does beyond the cleanup, E_{ext} . The premium is indexed by the fleet the action leaves behind. Replacing a tank therefore lowers next year's premium, and exit drives it to zero. The deductible is unchanged, because the facility already bears it in the payoff and the expected claim is net of it. The one term the model does not estimate is E_{ext} . I set it to the infant-health damage Marcus (2021) attributes to a leak, \$25,600 per leak in 2023 dollars, or 2.56 model units.³ Marcus counts infant health only and names what she omits, adult health, cancer, property values, and unreported leaks. The benchmark therefore prices a lower bound on the damage.

6.3 Measuring welfare

A contract changes welfare only through the choices it changes. Premiums and fund fees are transfers between a facility and its insurer, and they cancel in a social total. Net welfare counts the resources that do not cancel: the operating value facilities earn, the costs of changing tanks and opening stations, the cleanups, and the damage leaks do beyond the cleanup. I measure it as a present value per facility in the observed stock, in 2023 dollars, for the incumbents operating at the start and for the entrants that open afterward,

$$\begin{aligned} W &= W^I + W^E, \\ W^I &= PS + \text{Insurer} - \text{Gov} - \text{FirmAboveCap} - \text{External}, \\ W^E &= \text{Entry} + \text{Insurer}^E - \text{Gov}^E - \text{FirmAboveCap}^E - \text{External}^E. \end{aligned}$$

Every flow term is a discounted sum along the paths a policy induces. Let $\mu_g(s)$ be the share of observed facility-years in state g at state variables s , $\Pr_g(d \mid s)$ the choice probabilities under the policy, and M_g the transition matrix they induce, with exit absorbing. For a flow

³Marcus (2021) reports total damages over the leaks in her sample without stating a dollar base year. Mortality is 75.5 percent of the total and is valued at the EPA's value of a statistical life on its 2006-dollar convention. I therefore deflate the per-leak figure from 2006 to 2023 by a factor of 1.5114. Every other dollar in the premium is already in 2023 dollars.

$f_d(s)$ paid when a facility at s takes action d , with expected value $\bar{f}_g(s) = \sum_d \Pr_g(d | s) f_d(s)$, the present value over the observed stock is

$$\text{PV}_g(f) = \sum_s \mu_g(s) [(I - \beta M_g)^{-1} \bar{f}_g](s).$$

A facility that exits stops paying premiums and stops leaking, and its later flows are zero. The incumbent terms are

$$\begin{aligned} PS &= \frac{1}{\gamma} \sum_g \sum_s \mu_g(s) V_g(s), & \text{Insurer} &= \sum_g [\text{PV}_g(P_g) - \text{PV}_g(HS)], \\ \text{Gov} &= \frac{1}{\gamma} \sum_g \text{PV}_g(m \rho c^{\text{inst}}), & \text{FirmAboveCap} &= \bar{S}_{>L} \sum_g \text{Leaks}_g, \\ \text{External} &= E_{\text{ext}} \sum_g \text{Leaks}_g, & \bar{S}_{>L} &= E[(S^{\text{cleanup}} - \$1,000,000)^+]. \end{aligned}$$

Producer surplus is the facility's value function in dollars, averaged over the observed stock. Each of the other terms applies the present value to one flow. $\text{PV}_g(P_g)$ is the premium revenue the insurer collects, with flow $P_g(n')$, the premium on the fleet an action leaves in the ground. $\text{PV}_g(HS)$ is the claims it pays, with flow $H(n')S(n')$, the chance of a leak this year times the expected claim on it. The expected claim $S(n')$ is the fitted cleanup cost of a leak above the deductible and up to the \$1,000,000 coverage limit, evaluated at the post-action fleet. $\text{Leaks}_g = \text{PV}_g(H)$ uses the chance of a leak alone as the flow, and it counts the expected discounted number of leaks. The facility bears the expected cleanup above the limit, $\bar{S}_{>L}$, on each of those leaks. Government outlay has flow $m \rho c^{\text{inst}}$, the subsidized share of the install cost of the m tanks an action installs, and it is zero outside (CF3). External is the damage beyond the cleanup, charged at \$25,600 per leak under every contract, the baseline included, and every comparison therefore charges the same damage. Producer surplus and government outlay are divided by γ because they come out of the payoff. Premiums, claims, and damages are already dollars.

The entrant terms apply the same flows to the facilities that open. Let $\mathcal{E}_g$ be the discounted number of entrants in state g per facility in the observed stock, the sum over future years of the entry probability p_g times the pool remaining in that year, and let $\omega_g(s)$ be the distribution of state variables at which entrants open. An entrant flow has present value

$$\text{PV}_g^E(f) = \mathcal{E}_g \sum_s \omega_g(s) [(I - \beta M_g)^{-1} \bar{f}_g](s),$$

and Insurer^E , Gov^E , FirmAboveCap^E , and External^E are built from PV_g^E exactly as the incumbent terms are built from PV_g . The value of entry is the expected surplus of the entry decision itself. A potential entrant opens when $x_g + \varepsilon^E \geq \varepsilon^O$, where $x_g = \bar{V}_g - c^{\text{inst}} \bar{n}_g - ec$ is the value of a new facility net of the cost of opening it, and with logit shocks the choice is worth $\log(1 + e^{x_g})$ to each potential entrant over staying out. Summed over the discounted number of potential-entrant-years, $\mathcal{E}_g/p_g$,

$$\text{Entry} = \frac{1}{\gamma} \sum_g \frac{\mathcal{E}_g}{p_g} \log(1 + e^{x_g}) = \underbrace{\frac{1}{\gamma} \sum_g \mathcal{E}_g x_g}_{\text{mean value of the stations that open}} + \underbrace{\frac{1}{\gamma} \sum_g \frac{\mathcal{E}_g}{p_g} \mathcal{H}(p_g)}_{\text{entry-choice surplus}},$$

with $\mathcal{H}(p) = -p \log p - (1 - p) \log(1 - p)$. The first part books each entrant at the mean value of a new facility less the cost of opening it. The second is the private value that makes entry worthwhile for the stations that do open, which are the potential entrants with the best draws. Booking entry this way treats entrants as producer surplus treats incumbents, whose value includes their own shocks. The mean of the shocks is left out, because it belongs to the potential entrant’s outside alternative, which it has under every contract (Train 2009).

Section D.5 gives the closed form for $\mathcal{E}_g$ and the caveats on the severity input. The ledger has no consumer side, and the value of a new station to the drivers it serves is not counted. Section 5 notes that γ understates the response to a premium if facilities shop across carriers. In that case the dollar magnitudes below are overstated, and the rankings are more reliable than the levels.

6.4 Welfare results

Table 3 ranks the observed contracts and the four counterfactuals against the first best. Each is measured by net welfare, in present value per facility in the observed stock, counting incumbents and entrants and charging every contract the same infant-health damage per leak. I summarize each contract by the share of the first-best gain it achieves, its improvement over the observed contracts as a percent of the first best’s improvement, $100 (W - W^{\text{CF0}})/(W^{\text{FB}} - W^{\text{CF0}})$. The observed contracts score zero and the first best scores 100, and a contract that does worse than the observed contracts has a negative share. The first best raises net welfare from \$29,539 under the observed contracts to \$57,859 per facility. Table 4 reports what each contract does to leaks and to the age at which tanks are removed. Table 20 in Section D.6 reports the totals with larger pools of potential entrants.

Table 3: Net welfare ranked against the first best

Scenario	Net welfare	Change from (CF0)	Share of first-best gain (percent)	Rank
(FB) First best	57,859	28,320	100.0	—
(CF0) Observed contracts	29,539	—	0.0	2
(CF1) Flat fee in Texas	27,186	−2,353	−8.3	4
(CF2) Risk-rated pricing in the fund states	42,511	12,972	45.8	1
(CF3) Replacement subsidy	28,800	−739	−2.6	3
(CF4) Age-20 removal mandate	24,240	−5,299	−18.7	5

Notes: Present value of net welfare in 2023 dollars per facility in the observed stock, counting incumbents and entrants, with the pool of potential entrants at its baseline size. Each row holds the estimated parameters fixed, changes one policy, and resolves the facility’s problem and the entry rule, at an annual discount factor of 0.95. Net welfare is the identity of Section 6.3, with entrants booked at the expected surplus of the entry decision. Every row, the first best included, charges the same infant-health damage of 25,600 dollars per leak. The share of the first-best gain is the row’s change from (CF0) as a percent of the first best’s change from (CF0). Rank orders the observed contracts and the counterfactuals by net welfare, with 1 the highest.

Putting Texas on a flat fee (CF1) lowers net welfare by \$2,353 per facility relative to the

Table 4: Leaks avoided and the age at which tanks are removed

Scenario	Leaks avoided	SW band	DW band	SW gap to FB	DW gap to FB
(CF0) Observed contracts	—	6.6	5.2	0.8	0.8
(CF1) Flat fee in Texas	-0.006	6.6	5.2	0.8	0.8
(CF2) Risk-rated pricing in the fund states	0.039	6.3	4.9	0.5	0.5
(CF3) Replacement subsidy	0.000	6.6	5.1	0.8	0.6
(CF4) Age-20 removal mandate	0.142	5.6	5.0	-0.2	0.5
(FB) First best	0.129	5.8	4.5	—	—

Notes: Leaks avoided is the fall from (CF0) in the expected discounted number of leaks per facility in the observed stock, at incumbents and entrants. SW band and DW band are the mean age band at which incumbents' single-walled and double-walled tanks are removed, through downsizing, replacement, or exit, weighted by tank over the discounted path from the observed stock. Band 1 is ages 0 to 4, each higher band covers the next five years, and band 8 is 35 years and older, so a band of b corresponds to about $5b - 2.5$ years. The gap to FB is the row's mean band less the first best's, in bands. A negative gap means tanks come out younger than under the first best.

observed contracts, a share of -8.3 percent of the first-best gain. It adds 0.006 expected leaks per facility. Single-walled tanks come out in age band 6.6, about 31 years, 0.8 bands later than under the first best. The change in net welfare is -\$22,083 per Texas station, the stations whose contract changes, and +\$0 per control station, through entry alone.

Putting the fund states on risk-rated pricing (CF2) raises net welfare by \$12,972 per facility relative to the observed contracts, a share of 45.8 percent of the first-best gain. It avoids 0.039 expected leaks per facility. Single-walled tanks come out in age band 6.3, about 29 years, 0.5 bands later than under the first best. The change in net welfare is +\$14,519 per control station, the stations whose contract changes, and +\$0 per Texas station, through entry alone.

The replacement subsidy (CF3) lowers net welfare by \$739 per facility relative to the observed contracts, a share of -2.6 percent of the first-best gain. It changes what facilities do. The annual replacement rate moves from 0.39 to 0.82 percent of facilities and the exit rate from 1.58 to 1.49 percent. It avoids essentially no leaks. A subsidy pays for replacements that would have happened without it, and it keeps open sites that would have closed. The change in net welfare is -\$770 per Texas station and -\$736 per control station.

The removal mandate (CF4) lowers net welfare by \$5,299 per facility relative to the observed contracts, a share of -18.7 percent of the first-best gain. In its first year 55.6 percent of facilities exit and 58.9 percent of tanks close, against 1.6 and 1.9 percent under the observed contracts. It avoids 0.142 expected leaks per facility, and single-walled tanks come out in age band 5.6, 0.2 bands earlier than under the first best. A mandate retires sound tanks along with risky ones. A price lets each facility choose the margin it adjusts, and the mandate chooses one margin for every facility. The change in net welfare is -\$75,045 per Texas station and +\$3,019 per control station.

Under the first best the annual facility exit rate rises from 1.6 to 12.9 percent, the annual

tank closure rate rises by a factor of 6.7, and expected discounted leaks per incumbent facility fall by 61 percent. Single-walled tanks come out in age band 5.8 and double-walled tanks in band 4.5, the removal ages the other contracts are measured against. Entry nearly stops, with the discounted number of entrants per facility falling from 0.076 to 0.010. Because the premium collects the above-cap loss and the external damage, the insurer's position equals those two terms and incumbent net welfare equals producer surplus (Section D.5). The exit and closure rates are one-year rates on the observed stock, before the stock adjusts to the premium. The reform Section 4 measures raised tank closures by 1.3 percentage points a year on a control base of 1.92 percent. Set against the first best, the Texas repricing moved a small fraction of what full internalization would move, because the filed schedule prices the cleanup the insurer expects to pay and not the loss above the cap or the damage beyond it.

7 Conclusion

A compulsory price of coverage that carries no information about risk holds the riskiest capital in service. This paper measured that retention, and the effect of ending it, using Texas's 1998 fund closure, which moved 16,625 facilities from a flat per-tank fee to risk-rated private insurance at once. Closures of observably identical tanks rose 1.3 points a year (s.e. 0.3) on a 1.92 percent control base, and the rise was steepest among the oldest tanks, which the filed schedules charge most, at 3.9 points a year for tanks 20 years or older at the reform. Small and old facilities exited, and large facilities downsized and replaced instead.

The structural model, estimated on the filed prices and on the choices facilities make under the two contracts, compares pricing with the instruments states use to retire old tanks. I measure each against a first-best premium that charges every facility the full social cost of its leak risk, counting incumbents and entrants. Putting the fund states on risk-rated premiums raises net welfare by \$12,972 per facility, 45.8 percent of the gain the first best achieves over the observed contracts. It is the only instrument that improves on them. Returning Texas to a flat fee costs 8.3 percent of that gain, a replacement subsidy of 50 percent of the install cost 2.6 percent, and an age-twenty removal mandate for single-walled tanks 18.7 percent. A subsidy pays for replacements that would have happened without it and keeps open sites that would have closed. A mandate retires sound tanks along with risky ones, because it sets one margin for every facility, while a price lets each facility choose its own.

Risk rating closes less than half of the gap to the first best, because the filed schedules price the cleanup the insurer expects to pay and not the loss above the coverage limit or the health damage beyond the cleanup. The Texas repricing moved closures a small fraction of what full internalization would move.

Five boundaries limit how far these results carry. Texas is one treated state, and the inference rests on a single treated cluster and its 23 control states rather than on repeated treatment. The setting operates on the extensive margin, because an installed tank's wall cannot be changed, and the elasticities here need not transfer to hazards where care is continuous. The model prices each reform holding filed prices fixed and letting facilities enter without responding to one another, rather than the equilibrium a national repricing would produce. Entry carries part of the ranking. Counting incumbents alone, risk rating in the fund states

raises net welfare by \$12,949 per facility, and a flat fee in Texas raises it by \$225. The flat fee’s loss comes from the entrants it draws in, whom the ledger values at their private return with no benefit to the drivers they serve. And the dollar levels lean on measured inputs: a fitted mean severity transferred from other states’ claims, and a coverage-cost weight that understates the response to a premium when facilities shop across carriers, whose division makes the levels high-side even where the ranking is not.

The choice Texas made in 1998 is the choice the remaining fund states face now. Of the 36 state funds the EPA counts as financial-responsibility mechanisms today, six carry a statutory sunset, a claim cutoff, or a fee repeal dated between 2028 and 2030 (EPA 2026b).⁴ A fund that closes repeats the Texas decision with better information than Texas had. For those states the operative finding is the ordering. Pricing risk retires the riskiest capital at a lower cost than commanding its retirement, and a fund that keeps its uniform fee keeps the retention that fee finances. The result extends beyond tanks. Wherever compulsory coverage is met with uniform pricing, the same retention operates, and the price of coverage is a policy instrument whether or not a state uses it as one.

References

- Adda, Jérôme, and Russell Cooper. 2000. “Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies.” *Journal of Political Economy* 108 (4): 778–806.
- Aguirregabiria, Victor, and Pedro Mira. 2007. “Sequential Estimation of Dynamic Discrete Games.” *Econometrica* 75 (1): 1–53. <https://doi.org/10.1111/j.1468-0262.2007.00731.x>.
- Anderson, Soren T., and James M. Sallee. 2011. “Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards.” *American Economic Review* 101 (4): 1375–1409. <https://doi.org/10.1257/aer.101.4.1375>.
- Association of State and Territorial Solid Waste Management Officials. 2017. “Abandoned Underground Storage Tanks Report.” ASTSWMO Tanks Subcommittee. <https://astswmo.org/files/Resources/Tanks/2017-12-ASTSWMO-AbandonedTanksReport-FINAL.pdf>.
- ASTSWMO. 1999. “Annual State Fund Survey Archives.” Annual survey series, Association of State and Territorial Solid Waste Management Officials. <https://astswmo.org/annual-state-fund-survey-results/>.
- Barahona, Nano, Francisco A Gallego, and Juan-Pablo Montero. 2020. “Vintage-Specific Driving Restrictions.” *The Review of Economic Studies* 87 (4): 1646–82. <https://doi.org/10.1093/restud/rdz031>.
- Boomhower, Judson. 2019. “Drilling Like There’s No Tomorrow: Bankruptcy, Insurance, and Environmental Risk.” *American Economic Review* 109 (2): 391–426. <https://doi.org/10.1257/aer.20160346>.
- Boomhower, Judson, Meredith Fowlie, Jacob Gellman, and Andrew Plantinga. 2024. “How Are Insurance Markets Adapting to Climate Change? Risk Classification and Pricing in Homeowners Insurance.” Working Paper 32625. National Bureau of Economic Research.

⁴The six are Nebraska, Utah, Vermont, Illinois, Washington, and Missouri. Legislatures have extended such dates before, and each of these can be extended as well.

- <https://doi.org/10.3386/w32625>.
- Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney. 2012. “Pricing and Welfare in Health Plan Choice.” *American Economic Review* 102 (7): 3214–48.
- California State Water Resources Control Board. 2025. “Single-Walled Underground Storage Tanks: Timeline and Frequently Asked Questions.” Replacing Underground Storage Tanks (RUST) program history and single-walled closure timeline. <https://waterboards.ca.gov/ust/single-walled-tanks-progress/timeline-faqs.html>.
- Davis, Lucas, and Paige Weber. 2025. “Does Regulation Distort Exit Decisions? Evidence from u.s. Power Plants.” Working Paper 34454. National Bureau of Economic Research. <https://doi.org/10.3386/w34454>.
- Davis, Rebecca J., J. Scott Holladay, and Charles Sims. 2022. “Coal-Fired Power Plant Retirements in the United States.” *Environmental and Energy Policy and the Economy* 3: 4–36. <https://doi.org/10.1086/717217>.
- De Groote, Olivier, and Frank Verboven. 2019. “Subsidies and Time Discounting in New Technology Adoption: Evidence from Solar Photovoltaic Systems.” *American Economic Review* 109 (6): 2137–72. <https://doi.org/10.1257/aer.20161343>.
- Deaner, Ben, and Hyejin Ku. 2024. “Causal Duration Analysis with Diff-in-Diff.” <https://arxiv.org/abs/2405.05220>.
- Deshmukh, Ranjit, Paige Weber, Olivier Deschenes, Danae Hernandez-Cortes, Tia Kordell, Ruiwen Lee, Christopher Malloy, et al. 2023. “Equitable Low-Carbon Transition Pathways for California’s Oil Extraction.” *Nature Energy* 8 (6): 597–609. <https://doi.org/10.1038/s41560-023-01259-y>.
- Einav, Liran, Amy Finkelstein, and Mark Cullen. 2010. “Estimating Welfare in Insurance Markets Using Variation in Prices.” *Quarterly Journal of Economics* 125 (3): 877–921. <https://doi.org/10.1162/qjec.2010.125.3.877>.
- Einav, Liran, Amy Finkelstein, and Neale Mahoney. 2021. “The IO of Selection Markets.” *NBER Working Papers*, July. <https://ideas.repec.org/p/nbr/nberwo/29039.html>.
- EPA. 1988. “EPA Sets Standards for Underground Storage Tanks.” Press release, September 1988. <https://www.epa.gov/archive/epa/aboutepa/epa-sets-standards-underground-storage-tanks.html>.
- . 2014. “Energy Policy Act of 2005 and Underground Storage Tanks.” Other Policies and Guidance. <https://www.epa.gov/ust/energy-policy-act-2005-and-underground-storage-tanks>.
- . 2021. “UST Finder: National Underground Storage Tank and Release Data.” Feature service US_Underground_Storage_Tank_2019, ArcGIS Online item 88d551abd342485582c5ca4aac6ac0d6. <https://www.epa.gov/ust/ust-finder>.
- . 2024. “Underground Storage Tank Program Facts.” *Underground Storage Tank Program Facts*, May, 2. <https://www.epa.gov/system/files/documents/2024-05/ust-programfacts-may2024.pdf>.
- . 2026a. “Semiannual Report of UST Performance Measures, Mid Fiscal Year 2026 (October 1, 2025 – March 31, 2026).” EPA Office of Underground Storage Tanks. <https://www.epa.gov/system/files/documents/2026-06/performance-measures-mid-year-fy-26.pdf>.
- . 2026b. “State Financial Assurance Funds.” EPA web page, last updated April 29, 2026. <https://www.epa.gov/ust/state-financial-assurance-funds>.

- Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan. 2016. “Market-Based Emissions Regulation and Industry Dynamics.” *Journal of Political Economy* 124 (1): 249–302.
- Gruenspecht, Howard K. 1982. “Differentiated Regulation: The Case of Auto Emissions Standards.” *The American Economic Review* 72 (2): 328–31. <https://www.jstor.org/stable/1802352>.
- Hendren, Nathaniel. 2013. “Private Information and Insurance Rejections.” *Econometrica* 81 (5): 1713–62. <https://ideas.repec.org/a/ecm/emetrp/v81y2013i5p1713-1762.html>.
- House Research Organization. 1995. “Bill Analysis: HB 2587, Prohibiting Petroleum Tank Fund Reimbursements After 1998.” Texas House of Representatives.
- Indiana Capital Chronicle. 2026. “Excess Liability Trust Fund Coverage of Underground Storage Tank Replacement.” July 23, 2026.
- IRS. 2008. “Sole Proprietorship Returns, Tax Year 2006, Table 1: Nonfarm Sole Proprietorships: Business Receipts, Selected Deductions, Payroll, and Net Income, by Industrial Sectors.” Internal Revenue Service, Statistics of Income Division. <https://www.irs.gov/pub/irs-soi/06sp01br.xls>.
- Jacobsen, Mark R., James M. Sallee, Joseph S. Shapiro, and Arthur A. van Benthem. 2023. “Regulating Untaxable Externalities: Are Vehicle Air Pollution Standards Effective and Efficient?” *The Quarterly Journal of Economics* 138 (3): 1907–76. <https://doi.org/10.1093/qje/qjad016>.
- Jacobsen, Mark R., and Arthur A. van Benthem. 2015. “Vehicle Scrappage and Gasoline Policy.” *American Economic Review* 105 (3): 1312–38. <https://doi.org/10.1257/aer.20130935>.
- Kansas Department of Health and Environment. 2026. “UST Property Redevelopment Fund: Single-Wall System Replacement.” Kansas Storage Tank Act as amended by HB 2192 (2015). <https://www.kdhe.ks.gov/1045/UST-Property-Redevelopment-Fund---Single>.
- Kellogg, Ryan. 2014. “The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling.” *American Economic Review* 104 (6): 1698–1734.
- Kentucky Energy and Environment Cabinet. 2011. “UST System General Operating Requirements.” 401 KAR 42:030, Section 5(2). <https://eec.ky.gov/Environmental-Protection/Waste/underground-storage-tank/2011%20Regs/UST%20system%20general%20operating%20requirements.pdf>.
- . 2024. “Petroleum Storage Tank Environmental Assurance Fund Reimbursement Rates, 401 KAR 42:250.” Frankfort, KY. <https://eec.ky.gov/Environmental-Protection/Waste/underground-storage-tank/>.
- Louisiana Department of Environmental Quality. 2008. “Underground Storage Tank Trust Fund: Single-Wall to Double-Wall Upgrades.” LAC 33:XI, Chapter 14; La. R.S. 30:2195(F).
- Marcus, Michelle. 2021. “Going Beneath the Surface: Petroleum Pollution, Regulation, and Health.” *American Economic Journal: Applied Economics* 13 (1): 72–104. <https://doi.org/10.1257/app.20190130>.
- Massachusetts Department of Environmental Protection. 2017. “Underground Storage Tank Systems: Single-Walled Steel Tanks.” 310 CMR 80.15. <https://www.law.cornell.edu/regulations/massachusetts/310-CMR-80-15>.
- Mian, Atif, and Amir Sufi. 2012. “The Effects of Fiscal Stimulus: Evidence from the 2009 Cash for Clunkers Program.” *Quarterly Journal of Economics* 127 (3): 1107–42.

- <https://doi.org/10.1093/qje/qjs024>.
- Muehlenbachs, Lucija. 2015. “A Dynamic Model of Cleanup: Estimating Sunk Costs in Oil and Gas Production.” *International Economic Review* 56 (1): 155–85. <https://doi.org/10.1111/iere.12098>.
- Ringleb, Al H., and Steven N. Wiggins. 1990. “Liability and Large-Scale, Long-Term Hazards.” *Journal of Political Economy* 98 (3): 574–95.
- Rust, John. 1987. “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.” *Econometrica* 55 (5): 999–1033. <https://doi.org/10.2307/1911259>.
- Ryan, Stephen P. 2012. “The Costs of Environmental Regulation in a Concentrated Industry.” *Econometrica* 80 (3): 1019–61.
- Santos Silva, J. M. C., and Silvana Tenreyro. 2006. “The Log of Gravity.” *The Review of Economics and Statistics* 88 (4): 641–58. <https://doi.org/10.1162/rest.88.4.641>.
- Shavell, Steven. 1986. “The Judgment Proof Problem.” *International Review of Law and Economics* 6 (1): 45–58. [https://doi.org/10.1016/0144-8188\(86\)90038-4](https://doi.org/10.1016/0144-8188(86)90038-4).
- State of Arizona. 2019. “Noncorrective Action Tank Site Improvement.” A.R.S. 49-1071, enacted by Laws 2015, ch. 247. <https://www.azleg.gov/arsDetail/?title=49>.
- State of California. 2014. “Permanent Closure of Single-Walled Underground Storage Tanks.” Cal. Health and Safety Code 25292.05, added by Stats. 2014, ch. 547 (SB 445). https://leginfo.ca.gov/faces/codes_displaySection.xhtml?lawCode=HSC§ionNum=25292.05.
- . 2026. “Loans for Replacing, Removing, or Upgrading Underground Storage Tanks.” Cal. Health and Safety Code 25299.100 et seq. https://leginfo.ca.gov/faces/codes_displayText.xhtml?lawCode=HSC&division=20.&title=&part=&chapter=6.76.&article=.
- State of Maine. 2026. “Removal Schedule for Underground Oil Storage Tanks.” 38 M.R.S. 563-A(1). <https://legislature.maine.gov/statutes/38/title38sec563-A.pdf>.
- Stavins, Robert N. 2006. “Vintage-Differentiated Environmental Regulation.” *Stanford Environmental Law Journal* 25 (1): 29–63. <https://doi.org/10.2139/ssrn.848326>.
- Sweeting, Andrew. 2013. “Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry.” *Econometrica* 81 (5): 1763–803. <https://doi.org/10.3982/ECTA7473>.
- Texas Legislature. 1995. “House Bill 2587, 74th Legislature, Regular Session: Relating to the Regulation of Certain Aboveground and Underground Storage Tanks.”
- Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation*. 2nd ed. Cambridge: Cambridge University Press.
- U.S. Census Bureau. 2008. “County Business Patterns: 2006.” State file, released July 2008. <https://www.census.gov/data/datasets/2006/econ/cbp/2006-cbp.html>.
- . 2024. “Annual Retail Trade Survey: 2022.” Sales, gross margin, and total operating expense tables, released January 29, 2024. <https://www.census.gov/data/tables/2022/econ/arts/annual-report.html>.
- U.S. Department of Justice and Federal Trade Commission. 2023. “Merger Guidelines.” U.S. Department of Justice and Federal Trade Commission. <https://www.justice.gov/atr/2023-merger-guidelines>.
- U.S. General Accounting Office. 1992. “Hazardous Waste: Upgrading of Underground Storage Tanks Can Be Improved to Avoid Costly Cleanups.” Report to the Chairman,

- Environment, Energy and Natural Resources Subcommittee, Committee on Government Operations, House of Representatives GAO/NSIAD-92-117. U.S. General Accounting Office. <https://www.gao.gov/assets/nsiad-92-117.pdf>.
- UST Contractors. 2026. “Underground Storage Tank Removal and Replacement Cost Guide.” Trade cost guide. <https://ustcontractors.com>.
- Yin, Haitao, Howard C. Kunreuther, and Matthew W. White. 2007. “Do Environmental Regulations Cause Firms to Exit the Market? Evidence from Underground Storage Tank (UST) Regulations.” *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.1265611>.
- Yin, Haitao, Howard Kunreuther, and Matthew W. White. 2011. “Risk-Based Pricing and Risk-Reducing Effort: Does the Private Insurance Market Reduce Environmental Accidents?” *Journal of Law and Economics* 54 (2): 325–63. <https://ideas.repec.org/a/ucp/jlawec/doi10.1086-655804.html>.
- Yin, Haitao, Alex Pfaff, and Howard Kunreuther. 2011. “Can Environmental Insurance Succeed Where Other Strategies Fail? The Case of Underground Storage Tanks.” *Risk Analysis* 31 (1): 12–24. <https://doi.org/10.1111/j.1539-6924.2010.01479.x>.

A Data Appendix

A.1 Sources

Tank registries. The panel begins from the EPA’s UST Finder, the compendium that harmonizes every state’s tank registry into one national schema. That compendium leaves some states missing fields the design needs. I rebuild the registry for seven states directly from the state agency’s own database rather than the compendium extract: Texas, Arkansas, Colorado, Louisiana, Maine, Oklahoma, and Tennessee. Every rebuilt registry still reports the five fields the design requires: install date, closure date, wall type, fuel, and capacity. Texas records one field none of the other six carry, the carrier writing each facility’s financial-responsibility coverage in each year. That field maps a facility’s tanks to the rate filing that priced them.

Rate filings. [Section A.6](#) documents the six carrier rate engines built for Texas, the licensing rule that bounds their coverage, and the imputation applied where no filing survives.

Premium coverage and imputation. I reconstruct each Texas facility’s premium from six filed rate manuals covering the largest carriers in the market. Those manuals price 70.9 percent of insured Texas facility-years from 2008 through 2020. Two limits bound that coverage. A carrier writing without a Texas license files no rate manual with the state. That carrier leaves no schedule for me to recover. A tail of small carriers writes the remainder, each holding well under one percent of facility-years. I impute a premium for every unpriced facility-year and flag the imputation rather than drop the facility.

Selecting the largest carriers determines what the premium variable measures. These carriers write most of the market. The schedules I recover therefore capture what a typical Texas facility paid to hold a tank one more year. The unpriced facilities do not form a random subset. They concentrate in the later years of the panel, and reconstructed coverage thins after 2014. [Section A.6](#) documents the licensing rule, the carrier counts, and the imputation.

Trust fund claims. I collect facility-level cleanup claims from state remediation trust funds, covering 6 states. [Appendix C](#) states which cost field each fund reports and the rule that harmonizes them across states.

Release records. Every release record carries a report date, a facility identifier, and a release identifier. [Section A.3](#) states the rule that classifies each release as standalone, revealed by a closure, or already known at the time of a closure.

Fund fees. For the control states I use the ASTSWMO annual State Fund Survey. The survey reports each state’s uniform per-tank fee as raw text rather than a parsed number. Texas’s pre-reform entry states a combined figure for both tank types, one dollar amount for above-ground tanks and another for underground tanks, and I take the amount tied to the underground-tank label rather than the first dollar figure the string reports.

EPA performance measures. Every state reports to the EPA twice a year the percentage of its recently inspected tank facilities found in compliance with the federal technical standards, under one standard through fiscal 2021 and a differently defined one from fiscal 2019. [Section B.11](#) builds the enforcement comparison from the complement of that percentage.

A.2 Sample construction

Three fields decide whether a state can enter the sample at all: a tank’s install date, its closure date, and its wall type. The two dates place the tank in event time. Wall type, together with age, defines the risk class the market prices. State registries do not report these three fields consistently. I therefore begin with the 42 states that operated a uniform per-tank fund and keep the 24 whose registries report all three. That screen leaves Texas and 23 control states. I then restrict those registries to incumbents, the facilities holding at least one active tank on December 23, 1998. The reform repriced this stock. Tanks installed later entered a market that was already risk-rated and carry no before-and-after. [Table 6](#) reports the sample remaining after each restriction.

[Table 5](#) reports the registry audit behind that restriction, state by state. Each row gives the state, the provenance group its registry entered the panel through, and the share of tank records missing three of the fields the tank-level design needs: the closure date, the install date, and the wall type. A state’s registry enters through one of two documented routes, a bespoke cleaner reading the state’s own database or the EPA compendium, or it sits outside both routes for the reasons [Section A.1](#) gives. Heavy missingness on any one of the three fields removes a state from the tank-level design regardless of which route its registry took.

Table 5: Registry field completeness by state

State	Group	Percent missing			State	Group	Percent missing		
		Closure	Install	Wall			Closure	Install	Wall
TX	Target (TX)	0.0	0.0	9.9	CT	Excluded (Treated)	26.2	4.7	81.1
AL	Control Tier 1 (Custom)	9.4	24.1	85.0	FL	Excluded (Treated)	100.0	28.8	100.0
AR	Control Tier 1 (Custom)	7.7	18.5	21.7	IA	Excluded (Treated)	0.0	10.4	1.9
CO	Control Tier 1 (Custom)	6.9	21.6	24.7	MD	Excluded (Treated)	15.5	18.1	0.0
GA	Control Tier 1 (Custom)	15.5	3.0	10.0	MI	Excluded (Treated)	0.2	15.9	10.8
KY	Control Tier 1 (Custom)	0.8	16.1	12.2	NJ	Excluded (Treated)	4.3	0.2	0.9
LA	Control Tier 1 (Custom)	0.0	0.1	33.7	WI	Excluded (Treated)	100.0	79.7	59.3
ME	Control Tier 1 (Custom)	0.6	0.0	0.2	WV	Excluded (Treated)	4.3	15.4	8.2
MN	Control Tier 1 (Custom)	4.9	22.8	44.0	DE	Excluded (Other)	1.1	31.1	7.6
MO	Control Tier 1 (Custom)	26.0	35.3	0.3	NE	Excluded (Other)	43.3	3.6	75.7
MS	Control Tier 1 (Custom)	0.5	5.3	1.2	OR	Excluded (Other)	2.4	18.6	100.0
NC	Control Tier 1 (Custom)	0.0	9.7	9.2	WA	Excluded (Other)	100.0	5.1	100.0
ND	Control Tier 1 (Custom)	1.2	6.7	0.0	AK	Remainder	100.0	15.0	100.0
NM	Control Tier 1 (Custom)	2.5	14.1	4.6	CA	Remainder		2.6	0.2
OH	Control Tier 1 (Custom)	0.8	38.2	29.6	DC	Remainder	24.9	30.0	100.0
OK	Control Tier 1 (Custom)	3.1	13.7	0.0	HI	Remainder	100.0	19.4	100.0
PA	Control Tier 1 (Custom)	0.0	18.1	59.5	IL	Remainder	6.4	48.4	100.0
TN	Control Tier 1 (Custom)	10.8	7.9	9.6	IN	Remainder	100.0	53.3	100.0
UT	Control Tier 1 (Custom)	0.1	11.8	14.9	NH	Remainder	100.0	1.7	100.0
VA	Control Tier 1 (Custom)	7.9	13.6	5.9	NV	Remainder	100.0	1.8	0.0
ID	Control Tier 2 (EPA Std)	0.8	2.2	0.0	NY	Remainder	100.0	30.0	100.0
KS	Control Tier 2 (EPA Std)	16.5	9.4	0.0	RI	Remainder	100.0	1.4	78.0
MA	Control Tier 2 (EPA Std)	5.7	8.2	45.7	SC	Remainder	100.0	100.0	100.0
MT	Control Tier 2 (EPA Std)	4.7	9.9	12.2	VT	Remainder	0.5	100.0	100.0
SD	Control Tier 2 (EPA Std)	80.0	9.1	0.0	WY	Remainder	100.0	13.0	100.0
AZ	Excluded (Treated)	0.3	15.4	37.9					

Notes: The three numeric columns are in percent. Closure is the percent of tanks recorded as closed that have no closure date, and it is blank for a registry that records no closed tank. Install is the percent of tank records with no install date, and Wall is the percent with a missing or unknown wall type. These are the three fields the tank-level design requires. All three are computed from the state’s whole harmonized registry rather than from the study sample. Group is the study-group label from the registry build. Control Tier 1 (Custom) marks a state whose registry is read from the state’s own database, and Control Tier 2 (EPA Std) marks a state whose registry comes from the standard extract of the Environmental Protection Agency (EPA) compendium. Target (TX) is Texas, Excluded (Treated) and Excluded (Other) mark states the build sets aside from the control tiers, and Remainder is every other state. Group is a provenance label and not a statement of whether a state enters the estimation sample, and [Table 6](#) gives the sample restrictions. States are sorted by group and then by postal code, and the list continues in the right-hand block. Source: the registry completeness audit, Code/Verification/state_design_coverage.R.

Table 6: Sample construction

Restriction	Facilities	Tanks	Facility-years	Dropped
Harmonized state registries, all tanks	782,321	2,391,348	—	—
States in the study sample (Texas and 23 control states)	402,673	1,211,503	—	1,179,845 (49.3%)
Tanks with a recorded installation date	350,077	1,058,489	—	153,014 (12.6%)
Tanks with a recorded wall type, fuel type and installation cohort	304,677	905,594	—	152,895 (14.4%)
Tank-years in the annual panel, 1985–2020	295,429	869,794	5,357,470	35,800 (4.0%)
Tank-years at facilities flagged as gas stations	155,389	490,120	3,086,985	379,674 (43.7%)
Drop tank-years that install and close in the same year	154,190	486,533	3,085,688	3,587 (0.7%)
Tanks installed before the 1998 reform	135,308	415,498	2,752,739	71,035 (14.6%)
Facilities holding at least one tank active at the reform	79,404	281,311	2,294,577	134,187 (32.3%)
Birth-cohort coarsened-exact match on build, vintage and state	78,483	277,087	2,267,900	4,224 (1.5%)
Panel years 1990–2018	78,483	266,610	1,882,228	10,477 (3.8%)
Estimation sample: facility-year difference-in-differences	78,092	266,610	1,870,571	0 (roll-up)

Notes: Each row applies one restriction to the row above. Dropped reports the tanks removed by that row and the share of the tanks in the row above. Facility-years are undefined before the registry is expanded into an annual panel and are shown as dashes in the first four rows. A blank closure date is treated as an open tank and right-censored at 2020. No tank is dropped for a missing closure date. States whose registries never record closures contribute no closure events. The wall-type row also removes tanks with an unrecorded fuel type or installation cohort. Wall type, fuel type, and installation cohort enter the sample definition together. The final row removes no tanks: it aggregates the matched tank roster to facility-years, the unit of the regression. The panel-years row keeps the years every estimate uses. A tank drops there when none of its panel years falls inside them. North Carolina’s registry stops recording closures around 2012. Its years end there. The gas-station row keeps the facilities flagged as gas stations; every row after it, the match and the estimation sample are computed within that subsample, and the match is rebuilt inside it.

A.3 Variable definitions

The closure decision and the portfolio decision happen at different levels, and the analysis carries a panel for each. The tank-year panel records every tank in every year it sits in the ground and measures closure at the finest resolution the registries support. The facility-year

panel aggregates those tanks into a portfolio and measures the decisions an operator makes for the site as a whole, including replacement, consolidation, and exit.

Closure. I code the tank-level outcome as an annual closure indicator. It takes the value one in the calendar year matching the tank’s recorded permanent-removal date and zero in every other year the tank appears in the panel. For tank i , let t_i^{close} be that year. The indicator is

$$\text{closure}_{it} = \mathbf{1}\{t = t_i^{\text{close}}\}.$$

A tank with no recorded closure date stays in the panel through 2020 with the indicator at zero throughout. I treat a tank still in service at the end of the panel as right-censored rather than as one that never closed. The facility-level indicator takes the value one in any year when at least one of the facility’s tanks records a closure. Let c_{it} be the number of tanks facility i removes in year t , and let K_{it} be its total capacity. The facility indicators are

$$\text{any_closure}_{it} = \mathbf{1}\{c_{it} \geq 1\}, \quad \text{cap_cut}_{it} = \mathbf{1}\{K_{it} < K_{i,t-1}\}.$$

That distinction matters because the average facility operates 3.53 tanks and can retire one while continuing to trade on the others.

Replacement. I then separate closures by what the facility does next. I classify a tank closure as a replacement when the facility installs at least one new tank on a date later than that closure. The rule compares each closure date against the latest install date recorded anywhere in the facility’s history. A tank closed in 2003 and followed by an installation in 2011 therefore counts as a replacement. Tanks closed in the same calendar year they were installed are excluded before the classification runs. A closure that fails the test counts as permanent.

Contraction margins. The remaining margins compare the facility’s tank count and total capacity against the previous year. A facility contracts in a given year when it does not exit that year, ends the year holding strictly fewer tanks than it held the year before, and held positive capacity in the previous year. Within contracting facility-years I separate three cases by the proportional change in capacity,

$$g_{it} = \frac{\Delta C_{it}}{C_{i,t-1}} = \frac{C_{it} - C_{i,t-1}}{C_{i,t-1}}.$$

A facility consolidates when capacity stays within five percent in absolute value, the case of three tanks giving way to one tank of comparable total volume. It downsizes when capacity falls by more than five percent. It reconfigures upward when capacity rises by more than five percent, the case of several small tanks giving way to one substantially larger one. The three cases are

$$\text{downsize} : g_{it} < -0.05, \quad \text{consolidate} : |g_{it}| \leq 0.05, \quad \text{reconfigure up} : g_{it} > 0.05.$$

Exit and censoring. Consolidation earns its own margin because carriers price per tank. Cutting the number of tanks cuts the premium even when the facility keeps every gallon of

capacity. A facility exits in the last year it appears in the panel, provided that year falls strictly before 2020. Let t_i^{last} be the last year facility i operates any tank. Exit is

$$\text{exit}_{it} = \mathbf{1}\{t = t_i^{\text{last}} \text{ and } t_i^{\text{last}} < 2020\}.$$

That proviso is the censoring rule. I code the last row of a facility still operating in the panel’s final year as zero rather than as an exit.

Releases. I attach the release records to the panel to track pollution outcomes. Each release enters with a report date and a facility identifier. A tank closure requires an excavation. The excavation doubles as an inspection and can surface contamination already in the ground. I therefore classify each release by the number of days between its report date and the nearest tank closure at the same facility. I treat a release reported on the day of a closure, or within the sixty days following it, as revealed by that closure. I treat a release reported in a facility-year with no closure as standalone. I treat a release reported more than one hundred eighty days before a closure as already known at the time of that closure rather than revealed by it. I construct the same three-way split at thirty, forty-five, and ninety days as a check on the window.

A.4 Pre-reform covariates

Every covariate in the heterogeneity analysis is measured at the reform and held there, because the reform itself changes local market structure. A competitor count measured after 1998 would register an outcome of the treatment, and conditioning on it would absorb part of the effect the design estimates. Fixing each covariate before the treatment date keeps it a pre-treatment trait.

I place each facility in its 2000 Census tract by matching its coordinates to the tract polygon containing it, and every tract measure comes from that tract. Four indicators describe the neighborhood. A facility is rural when more than half its tract’s population is rural, under the Census 2000 urban-rural classification. It is low-population when its tract’s population density, in people per square mile, sits in the bottom third of the sample. It is low-income when its tract’s median household income, in 1999 dollars, is below the sample median. It is high-poverty when its tract’s poverty rate is above the sample median. One indicator describes the local market. For each facility I count the other gasoline-retailing facilities whose coordinates fall within one mile, or 1,609 meters, restricting the count to those holding at least one active tank on December 23, 1998. A facility is in a thin market when that count is one or fewer, and a retailer facing at most one nearby competitor can pass a cost increase through without losing customers.


Every covariate here enters only as a split of the treatment effect, never as a control in the headline regression.

A.5 Age distributions at the reform

Figure 9 plots the age distribution of the active tank stock in Texas and in the control states at the reform date and at three later snapshots, for the matched estimation sample. The

two groups hold a comparable stock at the reform, and the post-reform divergence appears as old tanks leaving the Texas distribution rather than as new tanks entering it.

Figure 9: Age distribution of the active tank stock, Texas and control states, at four snapshots



Notes: Share of active tanks at each single year of age, with ages of 40 and above grouped in the last bin. Age is the calendar year less the install year. The panels show the stock in 1998, the year of the reform, and in 2005, 2010, and 2015. Shares sum to one within each group and panel. Texas is red and the control states are blue, and dashed vertical lines mark each group's mean age. The text in each panel gives each group's mean age in years, the percent of its tanks aged 20 or more, and the number of active tanks. The sample is the matched sample of Table 6, in Texas and the control states. A tank is not counted in the year it is removed.

A.6 Carrier rate engines

For Texas I price each facility from the carriers' filed rate manuals, collected from the Texas and Iowa electronic rate-filing systems and linked to facilities through the Texas Commission on Environmental Quality contract records that name each facility's carrier. Six carriers have recoverable filings: Mid-Continent, Tank Owners Members (TOMICS), Great American, Zurich American, ACE, and Commerce and Industry (AIG). Together they price 70.9 percent of insured Texas facility-years. The main gap is two excess-and-surplus-lines carriers, Ironshore Specialty and Colony Insurance. Neither makes an admitted-market rate filing, and no schedule can be reconstructed for either. A facility insured by one of these keeps

its contract record but not a reconstructed premium, and its premium is imputed as the priceable-market mean for its age, wall, and era cell and flagged as imputed.

Every manual levies the premium per tank and loads it on the tank’s age and its wall construction, and beyond that the six differ more than one description can carry. Three of them set the base rate on age directly. TOMICS sets its base on install vintage, keyed to the 2006 Energy Act cutoff, and carries age as an additive load on top. Mid-Continent and Great American hold the base flat and move age through a load or a factor. The double-wall credit is the one feature common to all six, though the mechanism varies: Zurich keeps separate single-walled and double-walled base tables, ACE bundles double-wall with clad steel under a single construction code, and Mid-Continent’s credit requires double-walled steel. A double-walled fiberglass tank earns nothing under that credit. Only two of the six price the gallons a tank holds. That scarcity is why cutting capacity within a tank is the one adjustment margin the schedules largely ignore.

Rebuilding a premium means fixing three contract terms the records do not report. The first is the coverage limit. I hold every facility at the federal reference of one million dollars per occurrence and in aggregate, and the premium is therefore the price of the mandated floor rather than of any higher limit a firm bought. The second is the deductible, and I hold every facility at the filing’s reference deductible of five thousand dollars. The third is the underwriter’s schedule credit or debit, the discretionary adjustment for firm-specific risk. I cannot observe it. I price each facility at the schedule-neutral rate and report that as the standard premium, and I report a minimum and a maximum by moving the schedule to the edges of its filed band. Those bands are not uniform. Five of the six carriers apply a band of forty percent in each direction. Great American files no schedule rating at all, and its minimum, standard, and maximum coincide. Every carrier applies a minimum policy premium as a floor, and those floors run from two hundred and twenty-five to five hundred and twenty-five dollars across carriers. Weighted across carriers and deflated to 2023 dollars the floor is \$563 per policy-year, or \$219 per tank-year split across a facility’s tanks, and it binds for 15.0 percent of Texas facilities.

The manuals are filed as countrywide rate pages adjusted to each state, and I scale a national age-and-wall schedule to each state’s level by Harwayne’s method (Boor 1996). I assign the rate era from the policy effective date. I read each rate load from the facility’s records: tank age, wall construction, corrosion protection, overfill protection, stored product, and prior release. I set any load the records do not report to the filing’s reference tier. Zurich files its manual in Iowa, and I apply it to Texas facilities as an Iowa-rate proxy with a neutral territory factor. Each engine is coded from the filed factor tables, with every table read twice and reconciled before it is encoded, and each carries a self-test that reproduces hand-computed premiums from those tables. One filing required a departure from its printed text. Great American prints its composition rule as a sum, but the same filing states that its rates derive from ACE’s program, ACE’s manual states that the adjustments are multiplicative and ACE’s own worked sheets confirm it, and the additive reading would leave Great American’s age and construction factors inert. I price Great American multiplicatively. The resulting cards are stated on the same five-year age grid as the model’s state space, and the carrier-by-year variation they preserve is the dispersion that identifies γ in [Section 5](#).


A.7 The stock over time

The exhibits below describe the estimation sample rather than the reform’s effect. The first pair fixes the stock at the reform date. [Figure 10](#) gives the distribution of active tanks per incumbent facility on the day the Texas rule took effect, and [Figure 11](#) gives the distribution of total storage capacity per facility on the same date. Both are drawn for Texas and the control states together. The two groups enter the panel holding comparable portfolios, the condition the matched design requires.

The next four follow the same quantities through the panel. [Figure 12](#) and [Figure 13](#) plot tanks per facility and capacity per facility by year, and the Texas and control series separate after the reform rather than before it. [Figure 14](#) plots the mean age of tanks at removal. That age falls in Texas after the reform, as facilities retire tanks earlier in their lives than the control states do. [Figure 15](#) plots the raw annual closure probability, the unadjusted counterpart to the event study of [Section B.8](#).


Four further exhibits belonging to this appendix are drawn elsewhere because the sections that argue from them own them. The cleanup-cost distribution is [Figure 20](#) in [Section B.2](#), and the premium exhibits [Figure 30](#), [Figure 29](#), and [Figure 31](#) are drawn in [Appendix C](#).

Figure 10: Tank count per incumbent facility at the reform date




Notes: Distribution of the number of active tanks per facility on December 23, 1998, for incumbent facilities in Texas and the control states. Incumbent facilities hold at least one active tank at the reform, and the sample is every incumbent facility before matching. Each bar is the share of a group’s facilities with that tank count, and the shares sum to one within each group. Counts of eight or more are grouped in the last bar. Texas is gold and the control states are blue, and the bars of the two groups overlap.

Figure 11: Total tank capacity per incumbent facility at the reform date




Notes: Distribution of total storage capacity per facility on December 23, 1998, in thousands of gallons. Curves are kernel density estimates and the vertical axis is density. Texas is gold and the control states are blue. The sample is the incumbent facilities of [Figure 10](#), before matching, omitting facilities with zero capacity or more than 60,000 gallons.

Figure 12: Tanks per facility by year, Texas versus control states




Notes: Mean number of active tanks per facility by year, 1990 to 2020, for Texas (red) and the 23 control states (blue). The unit of observation is the facility-year, and the sample is the unmatched annual facility-year panel of these states, entrants included. A tank is active in a year if it is in the ground at any point in the year, and tanks installed and removed in the same year are excluded. In 2019 and 2020, 10 control states record no closures, and North Carolina's registry stops recording closures around 2012. The dotted vertical line marks the 1998 reform. Shaded bands are 95 percent confidence intervals for the mean, computed with facility-years treated as independent. The vertical axis does not start at zero.

Figure 13: Median facility capacity by year, Texas versus control states




Notes: Median total tank capacity per facility, in gallons, by year, 1990 to 2020, for Texas (red) and the 23 control states (blue). The unit of observation is the facility-year, and the sample is the unmatched annual facility-year panel of these states, entrants included. A facility's capacity is the sum of the capacities of its tanks in the ground at any point in the year, excluding tanks installed and removed in the same year. The median is taken across the facilities operating in each year. In 2019 and 2020, 10 control states record no closures, and North Carolina's registry stops recording closures around 2012. The dotted vertical line marks the 1998 reform. Shaded bands are distribution-free 95 percent confidence intervals for the median, computed with facility-years treated as independent.

Figure 14: Mean age of tanks at removal by year, Texas versus control states



Notes: Mean age, in years, of the tanks removed in each year, by removal year, 1990 to 2020, for Texas (red) and the 23 control states (blue). Age is the removal year less the install year. The unit of observation is a tank removal, and the sample is the matched sample of [Table 6](#). In 2019 and 2020, 10 control states record no closures, and North Carolina's registry stops recording closures around 2012. The dotted vertical line marks the 1998 reform. Shaded bands are 95 percent confidence intervals for the mean, computed with removals treated as independent.

Figure 15: Share of facilities closing a tank by year, Texas versus control states



Notes: Percent of facilities that close at least one tank in the year, by year, 1990 to 2020, for Texas (red) and the 23 control states (blue). The unit of observation is the facility-year, the sample is the unmatched annual facility-year panel of these states, entrants included, and tanks installed and removed in the same year are not counted as closures. In 2019 and 2020, 10 control states record no closures, and North Carolina's registry stops recording closures around 2012. The dotted vertical line marks the 1998 reform, the year of the federal upgrade deadline in every state. Shaded bands are 95 percent confidence intervals for the proportion, from the normal approximation with facility-years treated as independent.


B Additional Results

This appendix holds the results the main text states but does not display. [Section B.1](#) reports the full battery of checks behind the risk model’s fit, cell by cell and facility by facility, and [Section B.2](#) draws the cleanup-cost distribution behind the pure premium. [Section B.3](#) compares Texas and control-state stations at the reform, and [Section B.4](#) discusses the threats to the comparison that its controls do not remove. [Section B.5](#) reports the closure estimate as each control is added, [Section B.6](#) the cumulative closures by install year, and [Section B.7](#) the coefficients of the age-group split. [Section B.8](#) collects the facility event studies by margin, and [Section B.9](#) reports the static release estimates. [Section B.10](#) reports the county entry analysis, and [Section B.11](#) reports the enforcement-intensity comparison raised in [Section B.4](#).

B.1 Risk prediction model

[Figure 16](#) plots the out-of-fold predicted release probability from the risk equation of [Section 3.1](#) by average tank age and wall construction, averaged within cells of three-year average-tank-age bins by wall construction, the cells of the observed rates quoted in [Section 3.1](#). The predictions recover the observed profile cell by cell: the double-walled line rises with age, the single-walled line sits above it at every age, and the installation spike in the youngest bin appears in both. The fit statistics quoted in the main text, the AUC and the share of releases in the highest-risk decile, are computed from the same out-of-fold predictions.

Figure 16: Predicted release probability by tank age and wall construction




Notes: Each point is the mean out-of-fold predicted probability that a facility records its first confirmed release in the year, from the lasso logit of the risk equation of [Section 3.1](#), within a cell of the facility’s average tank age in three-year bins and wall construction. Values are in probability units, and 0.01 is one percentage point. The sample is facility-years before the facility’s first confirmed release in 24 states, Texas from 1990 to 1998 and the control states from 1990 onward. The solid line is facilities holding at least one single-walled tank, and the dashed line is facilities holding none. Facilities in New Mexico, whose registry records no wall construction, are not drawn. No facility is scored by a model fit that saw it. Shaded bands are pointwise 95 percent confidence intervals, from the delta method applied to a logistic refit without the penalty, with standard errors clustered by facility.

This appendix reports the three checks behind that fit in full. Every predicted probability is cross-fitted, as in the main text, and no facility is scored by a model that saw it.

The first check is concentration. If the model carries signal, the facilities it flags as riskiest should hold a large share of the releases that actually occur. [Figure 17](#) ranks every facility by predicted risk, from highest to lowest, and plots the share of next year’s releases the screened facilities contain against the share of facilities screened. A model with no signal traces the diagonal, since screening any share of facilities at random catches that same share of releases. My curve rises well above it. The highest-risk tenth of facilities holds 55.6 percent of the next year’s releases, and the highest-risk fifth holds 67.4 percent.


Figure 17: Risk concentration in the highest-ranked facilities



Notes: Out-of-sample concentration of risk. Facility-years are ranked by out-of-fold predicted probability of a first confirmed release from the cross-validated risk model, from the highest to the lowest, and the curve plots the cumulative share of first confirmed releases captured against the cumulative share of facility-years screened, both in percent. The 45-degree line is the no-information benchmark, and the dotted lines mark the highest-risk tenth of facility-years. The sample is 2,444,454 facility-years before the facility's first confirmed release in 24 states, Texas from 1990 to 1998 and the control states from 1990 onward. No facility is scored by a model fit that saw it.

A model can rank facilities correctly and still misstate how risky they are. The second check, calibration, asks whether the predicted probabilities match the rates facilities actually realize. [Figure 18](#) sorts facility-years into ten bins by predicted score and plots each bin's average predicted probability against its observed release rate. A perfectly calibrated model places every point on the forty-five degree line, where predicted risk equals realized risk. My points sit close to it. In the riskiest bin the model predicts a 13.5 percent chance of a release, and those facilities leak 13.1 percent of the time.


Figure 18: Calibration of predicted first-release probability



Notes: Out-of-sample calibration of the risk model. Facility-years are grouped into ten bins at the deciles of the out-of-fold predicted probability of a first confirmed release, and each point plots the bin mean of predicted first-release probability against the observed release rate in that bin, both in percent. The numbers label the bins from the lowest to the highest predicted probability, and the dashed line is the 45-degree line. The sample is 2,444,454 facility-years before the facility's first confirmed release in 24 states, Texas from 1990 to 1998 and the control states from 1990 onward. No facility is scored by a model fit that saw it.

The third check is the risk map itself. The model should reproduce the release rate in each age-and-wall cell, not merely rank facilities along it. [Figure 19](#) plots, for each cell, the rate the model predicts against the rate actually observed. My cells fall close to the forty-five degree line. As a single summary of how well the model separates future leakers from non-leakers out of sample, the area under the ROC curve is 0.81, with a 95 percent interval of 0.811 to 0.815, where 0.5 is chance and 1.0 is perfect. Against a base rate of 2.37 percent the area under the precision-recall curve is 0.237, with an interval of 0.233 to 0.241. The model recovers the level of risk in each cell. A risk-based premium prices exactly that.

Figure 19: Predicted versus observed release probability, by age-and-wall cell



Notes: Out-of-sample fit of the risk model at the cell level. Each point is a cell of the facility’s average tank age in three-year bins and wall construction, plotting the mean out-of-fold predicted probability of a first confirmed release in the year against the raw release rate observed in that cell. Both axes are in probability units, and 0.01 is one percentage point. Single-walled marks facilities holding at least one single-walled tank, and double-walled marks facilities holding none. Each label gives the age bin, and larger points have more facility-years. The dashed line is the 45-degree line. The sample is facility-years before the facility’s first confirmed release, Texas from 1990 to 1998 and the control states from 1990 onward, excluding New Mexico, whose registry records no wall construction.

The youngest tanks carry a spike of their own. Tanks age zero to two hold 6.2 percent of the facility-years behind Figure 16 and 11.9 percent of its releases. I keep this band in the figure and in the fitted model, because a release found at installation is a real claim against the fund and belongs in the premium. It does not carry the age gradient, and no result in this appendix depends on it.


The three checks above, and the cell map of Figure 16 and Figure 19, require a known wall type. New Mexico’s registry does not record one. Its facility-years enter the discrimination and calibration checks and drop out only of the cell-level figures. The excluded facility-years carry a release rate close to the sample average, and the restriction does not select on risk.

B.2 The distribution of cleanup cost

Figure 20 plots the distribution of cleanup cost across the 12,168 claims paid by 6 state trust funds, in 2023 dollars. The median claim runs \$106,000 against a mean of \$348,000. The

distribution is right-skewed, a minority of releases accounts for most of the spending, and the mean is the object a pure premium is built on. [Appendix C](#) states how each state's cost field is harmonized.

Figure 20: The distribution of cleanup cost across trust-fund claims



Notes: Kernel density of the cleanup cost of each claim paid by a state trust fund, pooled across states, in 2023 dollars on a logarithmic horizontal axis. The sample is 12,168 claims with a positive cost from the trust funds of 6 states (Colorado, Louisiana, New Mexico, Pennsylvania, Tennessee, and Utah). Costs are deflated to 2023 dollars by the month the claim starts. The dashed vertical line marks the median and the dot-dash vertical line marks the mean. Sources: state trust-fund claim records, and the consumer price index for all urban consumers from the Federal Reserve Bank of St. Louis (series CPIAUCSL).

B.3 Baseline characteristics

Texas and control facilities operate at the same scale. The median facility runs 3 tanks in Texas against 3 in the controls and holds 26 thousand gallons against 22. Both groups carry the same stock profile at the reform: a mean tank age of 15.0 years in Texas against 14.4 in the controls, 0.71 of tanks installed before 1989, and 0.89 single-walled. [Table 7](#) reports the comparison in full for the 78,483 facilities in the estimation sample, 16,444 in Texas and 62,039 in the controls, across 1,870,571 facility-years.

Table 7: Baseline characteristics, estimation sample

	Texas	Control	Norm. diff.
Stations	16,444	62,039	
Tanks per station, mean (standard deviation)	3.53 (1.67)	3.53 (1.86)	0.00
Tanks per station, median [interquartile range]	3 [3–4]	3 [2–4]	0.00
Total capacity, median [interquartile range] (thousand gallons)	26 [18–36]	22 [12–32]	+0.19
Mean tank age at the reform (years)	15.0	14.4	0.06
Single-walled share	0.91	0.89	+0.06
Share installed before 1989	0.76	0.70	+0.17

Notes: Estimation sample: facilities holding at least one tank active at the December 23, 1998 reform, in Texas and the 23 control states. Facility characteristics are measured at the reform date. Cells show means with standard deviations in parentheses, medians with interquartile ranges in brackets, and shares for indicator variables. Total capacity is in thousands of gallons. Mean tank age is the mean across the 255,619 incumbent tanks with a recorded install year, and its normalized difference is computed across tanks. Every other row is computed across facilities. The last column is the Imbens and Rubin normalized difference, the Texas minus control gap over the pooled standard deviation, and 0.25 in absolute value is the conventional balance threshold.

B.4 Threats to identification

This appendix discusses the threats to the comparison in Equation 2 that its controls do not remove.

Four other forces move a tank’s closure date, and each meets a different element of the design. The federal deadline moves it, and the deadline indicators and the age-by-type effects absorb the response common to tanks of the same type and age. Fuel demand and station margins move it, and the facility effect absorbs a facility’s own level while the year effect absorbs national movement. Local competition moves it, and no fixed effect in Equation 2 absorbs a change in it after the reform. Access to capital moves it, and I hold no measure of it. Reverse causation does not arise because carriers file rates statewide with the Texas Department of Insurance and those rates apply to a class rather than to an individual facility.

States legislate above the federal floor, and some order existing tanks out of the ground on criteria that bind within the sample window. Maine required the permanent closure of tanks installed before 1986 that had been retrofitted to meet leak detection and spill requirements, effective in 2008. Kentucky required the permanent closure of steel tanks and piping lacking corrosion protection in 2012, and required interior-lined steel tanks either to add cathodic protection or to close in 2013. Both states are controls. These rules, and the other closure mandates, upgrade mandates, and removal subsidies the control states enacted from 1991 to 2018, enter Equation 2 through the policy indicators in M_{st} , each from its effective year. Table 8 shows that adding them raises the estimate, as rules that close control-state tanks should.

A second kind of state rule reaches the sample through a different channel. Every state in the sample now requires double-wall construction on new installations, on a date it set for itself, and Texas adopted that requirement on January 1, 2009. A construction standard binding new tanks cannot force an old one to close, and grandfathering delays retirement rather than hastening it. The standard does raise the cost of replacing a tank because the replacement must meet the current standard at roughly seventy percent more per tank. Replacement is a margin this paper measures separately, and the Texas date falls a decade after the reform and inside the event-study window. [Section 4.3](#) reports the replacement response.

Enforcement intensity is the form of this concern that regulatory logic cannot dispose of, because a state that inspects more often finds more occasions to close a tank. Nothing in [Equation 3](#) absorbs a change in inspection effort specific to one state in one year, and the direction of the bias follows the state that changes. A rise in Texas inspections after 1998 would raise Texas closures for a reason other than price and bias the estimate upward. A rise in a control state would bias it toward zero.

Two measures bound the concern without disposing of it. Since the Energy Policy Act of 2005, federal law has required every state to inspect each regulated tank at least once every three years, and the EPA reports for every state and fiscal year how often inspectors looked and what they found. Texas looked less often than every control state, 102.4 inspections per thousand active tanks a year against a control mean of 168.5 over fiscal 2008 to 2024, and it found fewer violations when it looked, 21.1 percent of inspected facilities out of compliance against 32.8 percent under the standard reported through fiscal 2021 and 15.2 against 42.7 percent under its successor, ranked 3 of 22 and 2 of 22 from the lowest rate upward. The closures a state's enforcement forces are the product of how much of its stock it inspects, how often an inspection finds a violation, and how often a violation ends in a closure. The first two factors are measured, and both run against the concern. For enforcement to have raised Texas closures relative to the controls in the years the measures cover, Texas would have had to close tanks per violation found at least 2.5 times as often as the control states under the first standard and 4.6 times under the second, and at least 1.5 and 2.8 times had it inspected as large a share of its stock as they did. What the measures cannot do is sign the bias. Both series begin after the reform, and the bias in [Equation 3](#) is the change in the Texas-minus-control enforcement gap at the reform rather than its later level. [Section B.11](#) reports the measures, the bound, and the assumptions it rests on.

Anticipation operates in the same direction. The sunset passed in 1995 and gave operators three and a half years of notice, and a closure that an operator advanced falls in the pre-period and reduces the estimated increase. That lead time also places $k = -1$ within the announcement window rather than before it, which makes the event study's base year an imperfect untreated benchmark.

The estimation sample admits only facilities holding an active tank on December 23, 1998, and it records no exit before the reform. That zero is mechanical rather than behavioral. A facility that closed in 1994 is absent from the sample, and the pre-reform period carries no information about exit. The unrestricted registry carries those exits. An annual exit series for 1990 through 1998, built without the incumbent restriction, is the evidence on whether Texas and the control states followed the same pre-reform exit path.

B.5 Estimates as controls are added

Table 8 reports Equation 3 as the controls of Equation 2 are added one at a time, at the tank level in panel A and at the facility level in panel B. Column (1) holds only facility and year fixed effects. Column (2) adds the federal deadline indicators, column (3) a fixed effect for each age bin, and column (4) replaces it with an age-bin-by-type effect. Column (5) adds the control-state policy indicators and is the specification of the main text. Column (6) replaces the facility effect with a type-by-install-year effect at the tank level and a facility-type effect at the facility level. The age controls lower the estimate, the policy indicators raise it, and the estimate is positive in every column.

Table 8: Reform effect on closure as controls are added

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Tank closure</i>						
Texas $\times$ Post	0.93 (0.30)	1.00 (0.32)	0.56 (0.28)	0.43 (0.27)	0.71 (0.24)	0.42 (0.15)
Tank-years	5,587,786	5,587,786	5,587,786	5,587,786	5,587,786	5,587,782
<i>Panel B. Facility closes at least one tank</i>						
Texas $\times$ Post	0.74 (0.32)	0.57 (0.32)	0.53 (0.30)	0.30 (0.23)	0.52 (0.20)	0.24 (0.14)
Facility-years	1,882,228	1,882,228	1,882,228	1,882,224	1,882,224	1,882,224
Facility fixed effects	Yes	Yes	Yes	Yes	Yes	No
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Federal deadline indicators	No	Yes	Yes	Yes	Yes	Yes
Age-bin fixed effects	No	No	Yes	No	No	No
Age-bin $\times$ type fixed effects	No	No	No	Yes	Yes	Yes
Control-state policy indicators	No	No	No	No	Yes	Yes
Type $\times$ install-year (tank) or facility-type fixed effects	No	No	No	No	No	Yes

Notes: Estimates of Equation 3 in percentage points. Panel A is the tank-year panel and the outcome is tank closure. Panel B is the facility-year panel and the outcome is the closure of at least one tank. Column (5) is the specification of Figure 5 and Table 1. At the facility level the age bin is that of the facility's oldest tank and the type crosses the oldest tank's type with the facility's number of tanks at the reform. Standard errors clustered by state are in parentheses.

B.6 Cumulative closures by install year


Table 9 reports the cumulative effects from Equation 4 for all tanks in column (1) and separately by install year in columns (2) to (6), and Figure 21 plots the coefficients for each sample. Tanks installed before 1985 account for most of the extra closures. Twenty years after the reform, the reform had closed 1,200 more tanks installed before 1975 (s.e. 450), 1,100 more installed from 1975 to 1979 (s.e. 500), and 1,800 more installed from 1980 to 1984 (s.e. 1,300). For tanks installed from 1985 to 1988 it closed 460 more within five years and 50 more after twenty. Tanks installed from 1975 to 1984 show a positive pre-period in Figure 21, with coefficients averaging 1.0 to 1.1 points from 1990 to 1996. The columns are estimated separately and do not sum to column (1).

Table 9: Extra share of the reform-date stock closed in Texas

	(1) All tanks	(2) ≤ 1974	(3) 1975–79	(4) 1980–84	(5) 1985–88	(6) 1989–98
Extra closed after 5 years (per 100 tanks)	4.4 (1.4)	5.0 (2.5)	5.1 (2.8)	6.0 (3.7)	3.8 (0.6)	1.3 (0.4)
Extra closed after 10 years (per 100 tanks)	6.7 (2.3)	10.6 (4.1)	9.1 (4.8)	9.9 (6.7)	3.7 (0.9)	1.7 (0.8)
Extra closed after 15 years (per 100 tanks)	7.7 (3.2)	16.1 (5.8)	13.7 (6.6)	11.4 (8.9)	1.1 (1.5)	1.1 (1.4)
Extra closed after 20 years (per 100 tanks)	9.5 (3.8)	19.1 (7.1)	18.0 (8.2)	15.4 (11.0)	0.4 (2.4)	1.1 (2.1)
Extra tanks closed after 5 years	2,140 (700)	320 (160)	310 (170)	690 (430)	460 (80)	160 (50)
Extra tanks closed after 20 years	4,640 (1,840)	1,210 (450)	1,100 (500)	1,780 (1,270)	50 (300)	140 (260)
Would have closed after 5 years (%)	12.2	25.1	18.6	12.8	11.9	3.9
Would have closed after 20 years (%)	33.2	44.6	38.0	33.1	39.9	21.7
Pre-reform mean of θ_τ , 1990–96	0.33	0.35	1.09	0.97	0.08	-0.03
Texas tanks at the reform	48,760	6,358	6,108	11,518	12,200	12,576
Group and year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Weights			Tanks at risk			

Notes: Each column is one estimate of Equation 4, on groups of tanks of the same wall construction, fuel, capacity class, install year, and state, for the tanks at gas stations open at the end of 1998 in Texas and the control states. Column (1) pools all install years and columns (2) to (6) restrict to the install years in the column head. Extra closed after k years is $S_k(e^{\Delta H_k} - 1)$, the share of the Texas stock closed by year k beyond what would have closed without the reform, where ΔH_k sums the post-reform coefficients θ_τ through year k and S_k is the share of the Texas stock still open. Would have closed is $1 - S_k e^{\Delta H_k}$, the counterfactual share closed. Extra tanks closed multiplies the extra share by the Texas stock. The pre-reform mean is the average of θ_τ over 1990 to 1996 relative to 1997, in percentage points of cumulative hazard per year. Standard errors in parentheses are clustered by state and computed by the delta method.

Figure 21: Event study of the cumulative closure hazard, by install year



Notes: Coefficients θ_τ from Equation 4 estimated separately for all tanks and for each install-year group of Table 9, in percentage points. 1997 is the omitted year. Vertical bars are 95 percent confidence intervals from standard errors clustered by state.

B.7 Age-group estimates

Table 10 reports the coefficients of Equation 5. Replacing the age groups with a linear term in the tank's age at the reform, the reform effect is 1.3 points at the mean age of 15 years (s.e. 0.2) and rises by 0.19 points for each additional year of age (s.e. 0.01).

Table 10: Closure response by tank age at the reform


	Annual tank closure
Texas $\times$ Post	0.24 (0.18)
Texas $\times$ Post $\times$ age 10 to 14	0.88 (0.20)
Texas $\times$ Post $\times$ age 15 to 21	1.53 (0.28)
Texas $\times$ Post $\times$ age 22 or older	3.32 (0.30)
Texas $\times$ age 10 to 14	10.49 (0.50)
Texas $\times$ age 15 to 21	9.48 (0.71)
Texas $\times$ age 22 or older	8.08 (1.06)
Tank-years	5,587,786

Notes: Difference-in-differences estimates of Equation 5, in percentage points. Age groups are 9 years or younger, 10 to 14, 15 to 21, and 22 or older at the December 23, 1998 reform. The youngest group is the reference, and each Texas $\times$ Post $\times$ group row is the difference from it. Sample and controls follow Figure 5. Standard errors clustered by state are in parentheses.

B.8 Facility event studies


Figure 22 estimates Equation 2 at the facility level with the controls of Table 1, on the probability that the facility closes at least one tank. Figure 23, Figure 24, and Figure 25 do the same for downsizing, replacement, and consolidation. The pre-period is flat for any closure. Texas facilities replaced and downsized somewhat more than control-state facilities before the reform, with pre-period coefficients averaging 0.25 and 0.16 points.

Figure 22: Event study of tank closure at the facility level, Texas relative to control states




Notes: The outcome is an indicator that the facility closes at least one tank in the year. Coefficients from Equation 2 estimated on the facility-year panel with the controls of Table 1, in percentage points. 1997 is the omitted year. Vertical bars are 95 percent confidence intervals from standard errors clustered by state.

Figure 23: Event study of facility downsizing, Texas relative to control states




Notes: The outcome is an indicator that the facility downsizes in the year. Coefficients from Equation 2 estimated on the facility-year panel with the controls of Table 1, in percentage points. 1997 is the omitted year. Vertical bars are 95 percent confidence intervals from standard errors clustered by state.

Figure 24: Event study of facility replacement, Texas relative to control states



Notes: The outcome is an indicator that the facility replaces a tank in the year. Coefficients from Equation 2 estimated on the facility-year panel with the controls of Table 1, in percentage points. 1997 is the omitted year. Vertical bars are 95 percent confidence intervals from standard errors clustered by state.

Figure 25: Event study of facility consolidation, Texas relative to control states



Notes: The outcome is an indicator that the facility consolidates in the year. Coefficients from Equation 2 estimated on the facility-year panel with the controls of Table 1, in percentage points. 1997 is the omitted year. Vertical bars are 95 percent confidence intervals from standard errors clustered by state.

B.9 Release estimates

Table 11 reports Equation 3 on each release outcome at the facility level, with the controls of Table 1.

Table 11: Release-discovery response, by outcome

	(1) Any release	(2) Outside a closure	(3) At a closure
Texas $\times$ Post	-0.82 (0.23)	-0.27 (0.21)	-0.55 (0.08)
Control mean, post-reform (%)	2.67	2.09	0.58
Facility-years	1,882,224	1,882,224	1,882,224

Notes: Facility-level estimates of Equation 3 in percentage points, one regression per column, with the controls of Table 1. Any release pools every confirmed release, outside a closure keeps releases discovered outside a closure inspection, and at a closure keeps releases reported within the sixty days following a closure. Standard errors clustered by state are in parentheses.

B.10 County entry, synthetic control

The reform had no perceptible effect on entry. The closure results describe incumbent facilities, and this section asks whether the rate at which new facilities open changed with the reform.

A facility enters in the year of its earliest tank installation. I sum entrants from the county panel to state-year counts, for facilities with no undated tank, and compare Texas with the 23 control states on four outcomes. These are entrants per 100,000 residents, the log of entrants, entrants as a percent of the previous year's stock of facilities, and the raw count. I estimate the effect by synthetic difference-in-differences and by synthetic control. Both choose weights on the control states to match Texas over 1990 to 1998. As a comparison I also use difference-in-differences with equal weights. I report the average effect over 1999 to 2006 and over 1999 to 2017. Inference is by placebo. Each control state is assigned the treatment in turn, with Texas kept as a donor, and the placebo p -value is one plus the number of control states whose estimate is at least as large in absolute value as the Texas estimate, divided by 24.

Table 12 reports the estimates, and Figure 26 plots Texas against its synthetic counterparts. On entrants per 100,000 residents the synthetic difference-in-differences, synthetic control, and difference-in-differences estimates for 1999 to 2006 are 12.3, 7.8, and 35.2 percent, with placebo p -values of 0.82, 0.68, and 0.41.


Ten of the twelve estimates have placebo p -values of 0.09 or higher in both windows. They are the three estimators on entrants per 100,000 residents and on the entry rate, and the synthetic difference-in-differences and difference-in-differences estimators on log entrants and raw entrants. The exceptions are the synthetic-control estimates for log entrants and raw entrants, with all of their weight on Pennsylvania and placebo p -values of 0.045 in both windows.

Table 12: Effect of the reform on entry of new facilities, by outcome and estimator

Outcome	Estimator	Effect 1999–2006	p 1999–2006	Effect 1999–2017	p 1999–2017	Pre-period RMSPE
Entrants per 100k	SDID	+12.3%	0.818	+12.5%	0.864	0.35
Entrants per 100k	SC	+7.8%	0.682	+3.0%	0.864	0.18
Entrants per 100k	DiD	+35.2%	0.409	+37.0%	0.455	0.65
Log entrants	SDID	+30.4%	0.409	+57.4%	0.273	1.83
Log entrants	SC	+165.5%	0.045	+260.8%	0.045	0.47
Log entrants	DiD	+47.6%	0.182	+85.4%	0.136	1.71
Entry rate (% of stock)	SDID	+38.9%	0.227	+48.3%	0.182	0.19
Entry rate (% of stock)	SC	+32.9%	0.318	+51.0%	0.091	0.13
Entry rate (% of stock)	DiD	+52.6%	0.136	+63.5%	0.091	0.32
Raw entrants	SDID	+20.6%	0.136	+19.5%	0.136	204.28
Raw entrants	SC	+58.6%	0.045	+57.6%	0.045	120.91
Raw entrants	DiD	+3.8%	0.727	-4.1%	0.727	242.38

Notes: Estimated effect of the reform on the entry of new facilities, Texas against the 23 control states, by synthetic difference-in-differences (SDID), synthetic control (SC), and difference-in-differences with equal weights (DiD). A facility enters in the year of its earliest tank installation. Entrants are summed from the county panel to state-year counts, for facilities with no undated tank, over 1990 to 2017. Entrants per 100k is entrants per 100,000 residents, log entrants is the natural log of the count, the entry rate is entrants as a percent of the previous year's stock of facilities, and raw entrants is the count. Effects are averages over 1999 to 2006 and over 1999 to 2017, in percent of the Texas mean over 1990 to 1998 (1991 to 1998 for the entry rate), and for log entrants as 100 times the exponential of the estimate minus one. The p -value is a placebo p -value. Each control state is assigned the treatment in turn, with Texas kept as a donor, and p is one plus the number of control states whose estimate is at least as large in absolute value as the Texas estimate, divided by 24. The smallest attainable value is 0.042. Pre-period RMSPE is the root mean squared prediction error of the fit to Texas over the pre-reform years, in the units of the outcome, for the 1999 to 2006 fit. Sources: state tank registries, and county population from the Bureau of Economic Analysis.

Figure 26: Texas entry and synthetic Texas, by outcome



Notes: Entry of new facilities in Texas and in synthetic Texas over 1990 to 2017, one panel per outcome. Synthetic Texas is a weighted average of the 23 control states. The panel labeled entrants per 100 incumbents is the entry rate, entrants as a percent of the previous year’s stock of facilities. The synthetic paths are the synthetic difference-in-differences and synthetic control fits for the 1999 to 2017 window, with weights chosen to match Texas over the pre-reform years. The shaded years 1995 to 1998 mark the announcement window, and the dashed line marks the reform. Sources: state tank registries, and county population from the Bureau of Economic Analysis.

A failure to reject is not an estimate of zero. The point estimates are positive in 23 of the 24 cells. With 23 control states, a rejection at the five percent level requires Texas to rank first among the 24 states. Entry is dated by the earliest tank installation in each state registry, and the headline sample leaves out facilities with an undated tank.

B.11 Enforcement intensity


Section B.4 raises one threat that the regulatory record cannot dispose of: Texas may have inspected its tanks harder after the reform than the control states inspected theirs, and closed tanks on that account rather than on the price. Two measures of enforcement exist for every state, how often inspectors looked and what they found. This section reports both, states what they show, and then says what the pair can and cannot tell the design. Both are level comparisons over the years after the reform. Neither is a test of a change at the reform, and nothing here is a treatment effect.

The measures. EPA publishes, for every state and fiscal year, the number of on-site inspections conducted under the Energy Policy Act and the percentage of the facilities inspected in the preceding twelve months that were found in compliance with the federal technical standards. Inspection coverage is the first count per thousand active tank systems, from fiscal 2008 through fiscal 2024. It is per tank rather than per facility, because the source publishes no state facility count. The failure rate is the complement of the second, the share of inspected facilities found out of compliance. Two standards make up the failure series. Significant operational compliance, SOC, runs from fiscal 2004 through fiscal 2021, and Texas first reports it in fiscal 2005. The technical compliance rate, TCR, replaces it from fiscal 2019 through fiscal 2024. The two are defined differently and need not share a base. States migrated from one to the other one at a time between fiscal 2019 and fiscal 2021, and in those years each standard rests on a changing subset of states. I never average, splice, or draw the two as one line.

Texas inspected less often and found less when it did. Over fiscal 2008 to 2024, Texas conducted 102.4 inspections per thousand active tanks a year, against 168.5 in the twenty-one control states of the closure design and 167.1 in all other states. Ranked from the fewest inspections per tank upward, Texas stands at 1 of 22 among the control states and itself and at 3 of 51 among the fifty states and the District of Columbia. Over the SOC years, 21.1 percent of inspected Texas facilities were found out of compliance, against 32.8 percent in the control states and 31.5 percent in all other states, gaps of -11.8 and -10.4 points. Over the TCR years the Texas rate is 15.2 percent against 42.7 and 42.7 percent, a gap of -27.5 points against either group. Ranked from the lowest failure rate upward, Texas stands at 3 of 22 under SOC and 2 of 22 under TCR among the control states and itself, and at 10 of 51 and 7 of 51 among the fifty states and the District of Columbia. [Figure 27](#) and [Figure 28](#) plot the two series by fiscal year. Texas reported few inspections in fiscal 2008 to 2010, and it sits below the control mean in every year of the coverage series. [Table 13](#) and [Table 14](#) report the gaps as regressions of the state-year measure on a Texas indicator, against all other states in Panel A and against the control states in Panel B. For the failure rate the rows are the pooled series and each standard alone, and the second column adds one effect for each standard in each fiscal year. On the pooled row those effects confine every comparison to states reporting the same standard in the same year, and the pooled gap can be read in that column alone. For coverage, with no split by standard, the second column adds fiscal-year effects. There are no state effects, because a Texas indicator that does not vary over time would be absorbed by them. Standard errors are clustered by state, and no significance test accompanies them, because one treated state and a descriptive comparison offer no sharp null. The ranks are rankings of realized state means over each panel's window, not statistics of the regressions. Both series begin after the December 1998 reform, and no reform line is drawn.

What the comparison can say. The closures a state's enforcement forces in a year are the product of three things: the share of its stock its inspectors reach, the share of the inspected facilities found in violation, and the share of violations that end in a closure. Per facility, $E_{st}/N_{st} = \rho_{st} y_{st} \pi_{st}$. The first two factors are the measures above, with coverage standing in for the inspected share per tank rather than per facility, and the third is observed nowhere. Both measured factors are lower in Texas. For the product to have been higher in Texas

Figure 27: Inspection coverage, Texas against the comparison states




Notes: On-site inspections conducted under the Energy Policy Act per 1,000 active tank systems, by fiscal year. The measure is per tank rather than per facility, because the source publishes no state facility count. The statute’s one inspection per facility every three years is therefore not drawn as a reference line. The series begins in fiscal 2008, the first year the inspection count is reported. The control states are the 23 states of the closure design, and all states are the forty-nine remaining states plus the District of Columbia. Both comparison lines are unweighted state averages and exclude Texas. Nothing in the figure is a treatment effect. Source: Environmental Protection Agency, underground storage tank performance measure reports.

than in the control states, Texas would have had to close tanks per violation found at least 2.5 times as often as the controls over the SOC years and 4.6 times over the TCR years, at the inspection rates each reported. Had Texas inspected as large a share of its stock as the control states, the thresholds would be 1.5 and 2.8. That pair is the more conservative one, because a few control states inspect far more often than the rest and pull the control mean up.

What the estimator picks up is not that product but its change. Equation 3 compares the change in Texas closures at the reform with the change in the control states, and enforcement biases it only if the Texas-minus-control enforcement gap moved at the reform. The failure series begins in fiscal 2005 and the coverage series in fiscal 2008, six and nine years after the reform, and the pre-reform gap is never observed. The comparison therefore bounds how large the enforcement channel can have been in the years it covers, and it cannot sign the bias. A Texas program that enforced far less than the controls before the reform and less so after it would bias the estimate upward while showing fewer inspections and fewer violations in every reported year. The years the bound does not reach include 1999, when the federal

Figure 28: Enforcement failure rate, Texas against the comparison states



Notes: Failure rate among the facilities inspected in the preceding twelve months, 100 minus the published combined compliance rate, by fiscal year, in percent. Federal law requires every facility to be inspected at least once every three years from 2007. Each year therefore covers about one third of the stock, and any three consecutive years cover every facility in the state. Significant operational compliance (SOC, solid lines) and the technical compliance rate (TCR, dashed lines) are different standards on possibly different denominators, drawn as separate series and never spliced. The shaded years, fiscal 2019 to 2021, are the reporting transition, when states migrated individually and each line rests on a changing subset of states. The control states are the 23 states of the closure design, and all states are the forty-nine remaining states plus the District of Columbia. Both comparison lines are unweighted state averages and exclude Texas. The series begins after the Texas reform, no reform line is drawn, and nothing in the figure is a treatment effect. Source: Environmental Protection Agency, underground storage tank performance measure reports.

upgrade deadline moved closures in every state. The coefficients for 2001 to 2004 in [Figure 5](#) are of the same size as those in the covered years.

Two features of the measures decide how far they describe the stock rather than the inspectors. The inspected set is fixed by statute rather than chosen. Since fiscal 2008 each year's inspections are the roughly one third of a state's facilities whose turn in the three-year cycle fell due, and three consecutive years reach every facility ([EPA 2014](#)). A year's failure rate departs from the stock's only through the order a state works its cycle, through re-inspection of facilities written up, and through repairs made between visits, and the last two push the measured rate above the stock's. The lower Texas failure rate then fits two readings, and neither is established here. Texas inspectors may have held facilities to the same standard

Table 13: Inspection coverage, Texas against the comparison states

	Raw	Year FE
<i>Panel A. Texas versus all states (50 states plus DC)</i>		
Inspection coverage	-64.76 (10.72)	-64.66 (10.80)
$N = 859$, states = 51, fiscal years = 17		
<i>Panel B. Texas versus the control states (21)</i>		
Inspection coverage	-66.11 (20.02)	-65.93 (20.42)
$N = 368$, states = 22, fiscal years = 17		
Control mean, excluding compartment reporters: 170.7 per 1,000 active tanks		
Year FE	No	Yes

Notes: State by fiscal year, fiscal 2008 to 2024. The outcome is on-site inspections per 1,000 active tank systems, counted per tank rather than per facility, because the source publishes no state facility count. Panel A compares Texas against all states, Panel B against the control states of the closure design. Cells report the coefficient on a Texas indicator, in inspections per 1,000 active tanks, with the standard error clustered by state in parentheses beneath. The second column adds fiscal-year effects. No state effects enter, because a time-invariant Texas indicator would be absorbed by them. Ranked from the fewest inspections per tank upward, Texas stands at 1 of 22 in Panel B and at 3 of 51 in Panel A. These are rankings of realized state means over the window and are not statistics of the regressions. The control mean excluding the states that report tank compartments rather than tanks is shown inside the table, because those states' inflated denominators understate their coverage. Nothing in the table is a treatment effect, and no significance test is reported. Source: Environmental Protection Agency, underground storage tank performance measure reports.

and met a sounder stock, though that stock is itself a survivor of the reform and its condition is not independent evidence about enforcement. Texas inspectors may instead have been less demanding, and laxity lowers the violations written up while raising the severity of those that are. The product the bias turns on is not signed by either reading. Fewer inspections per tank, the margin the threat turns on most directly, needs no such reading and runs against the threat on its own.

What remains open. The measures have no pre-reform observation, and the state's own enforcement records before 1999 are what would turn the bound into a measurement. Coverage is per tank rather than per facility. The statute's requirement is per facility, a state whose facilities hold more tanks shows lower per-tank coverage at the same per-facility frequency, and the source publishes no state facility count to convert one to the other. Thirteen states report tank compartments rather than tanks, several of them control states, and their inflated denominators understate their coverage. Excluding them raises the control mean to

Table 14: Enforcement failure rate, Texas against the comparison states

	Raw	Standard $\times$ year FE
<i>Panel A. Texas versus all states (50 states plus DC)</i>		
All measures (pooled SOC and TCR)	-13.87 (1.47)	-13.54 (1.48)
$N = 1,018$, states = 51, fiscal years = 21		
SOC	-10.41 (1.50)	-10.08 (1.51)
$N = 811$, states = 51, fiscal years = 18		
TCR	-27.48 (2.70)	-27.40 (2.70)
$N = 207$, states = 51, fiscal years = 6		
<i>Panel B. Texas versus the control states (21)</i>		
All measures (pooled SOC and TCR)	-14.82 (2.26)	-14.44 (2.32)
$N = 451$, states = 22, fiscal years = 21		
SOC	-11.75 (2.41)	-11.44 (2.44)
$N = 364$, states = 22, fiscal years = 18		
TCR	-27.47 (3.60)	-26.60 (4.02)
$N = 87$, states = 22, fiscal years = 6		
Standard $\times$ year FE	No	Yes

Notes: State by fiscal year. The outcome is the failure rate among facilities inspected in the preceding twelve months, 100 minus the published combined compliance rate, in percent. Panel A compares Texas against all states, Panel B against the control states of the closure design. Cells report the coefficient on a Texas indicator, in percentage points, with the standard error clustered by state in parentheses beneath. The second column adds one effect per standard and fiscal year, and on the pooled row that confines each comparison to states reporting the same standard in the same year. No state effects enter. SOC, significant operational compliance, spans fiscal 2004 to 2021 and TCR, the technical compliance rate, fiscal 2019 to 2024. The fiscal 2019 to 2021 TCR cells rest on few states. Ranks run from the lowest realized mean failure rate upward. A low rank is a low failure rate. Texas stands at 3 of 22 under SOC and 2 of 22 under TCR in Panel B, and at 10 of 51 and 7 of 51 in Panel A. The ranks are of realized state means over each panel's window, not statistics of the regressions. The series begins after the Texas reform, nothing in the table is a treatment effect, and no significance test is reported. Source: Environmental Protection Agency, underground storage tank performance measure reports.

170.7 and widens the gap. Texas reported few inspections in fiscal 2008 to 2010, the first three years of the coverage series, and its coverage from fiscal 2011 sits closer to the control mean while remaining below it in every year. The Texas SOC mean pools fiscal 2005 to 2007,

before the statutory cycle was complete. The state means are unweighted while the closure design weights tank-years within cells. The closures forced per violation are not observed.

C The Claims-Implied Premium

This appendix asks whether the premiums Texas carriers file track the expected cost of the risk they insure. The benchmark is the claims-implied premium, or fair premium, the expected annual cleanup cost of a facility’s leak risk. I build it first, set the filed premiums against it second, and state last what the comparison cannot show.

The cost object here is not the one [Section 3.1](#) quotes. That section prices the risk the structural model consumes, a national age-by-wall hazard multiplied by the claims-fitted severity of the cell, measured at the facility. The fair premium below differs on three counts. It uses each facility’s own fitted hazard in place of its cell average. It measures severity as the amount the fund pays after the operator’s deductible rather than the gross cost of the cleanup. It then splits the facility figure across tanks. The two objects answer different questions and need not agree in level.

C.1 The fair premium

The fair premium is the expected cost of a facility’s leak risk in a year, the product of how often it leaks and what a leak costs,

$$P_i^{\text{fair}} = \hat{h}_i \hat{S}_i.$$

The frequency $\hat{h}_i$ is the facility’s fitted first-release probability from [Section 3.1](#). The severity $\hat{S}_i$ is its fitted cleanup cost net of the deductible the facility bears, making it the share the state fund pays. Both are conditional means rather than realized amounts. Their product is what a break-even insurer would need to collect from facilities like this one, not what any single facility went on to cost. A facility charged less than its own expected cost has the remainder paid by someone else, in this case by the gasoline taxpayers who finance the fund and by the safer facilities in the pool. Both objects are measured at the facility, and the frequency has to be, because a release is recorded against the facility and is never attributable to a particular tank.

The severity needs its own data, because a cleanup cost is observed only when a facility files a claim. I collected trust-fund claims records from six states: Colorado, Louisiana, New Mexico, Pennsylvania, Tennessee, and Utah. All six enter the model fit, since the fit needs every cleanup it can get. Texas has no fund and files no claims. Its severity is taken out of sample as the average state effect across the fit states excluding Louisiana, an offset of -0.33 on the log scale. A claim is the amount a fund paid to clean up one confirmed release. Each claim links to the facility that filed it through the state’s site registration number, attaching to the claim the same traits the hazard model uses: tank age, wall construction, fuel, and capacity. The result is 12,168 claims, each one cleanup paired with the traits of the facility it came from.

Harmonizing the cost field takes one decision per state. I measure the fund-payable cleanup cost per release, net of the facility’s deductible, deflated to 2023 dollars by the CPI-U, BLS series CPIAUCSL, rebased to the 2023 annual average. Colorado, New Mexico, and Tennessee report the closed, realized cost of completed cleanups. Colorado’s field is the amount reimbursed to the operator, and Tennessee’s is the paid amount on closed claims only. Pennsylvania reports the ultimate incurred loss. That measure is paid cost plus case reserves, and it retains large claims while they remain open. The state’s most expensive releases are therefore not truncated. The states also record the operator’s share differently, and I take each at the point where its own records already net it out rather than imposing one rule on all. Each of the three above is fund-paid by construction. New Mexico is the one state whose recorded amount is gross of the operator’s share, and I subtract its \$10,000 deductible directly. Louisiana reports a recommended amount across applications, and I divide by the application count to reach one release. Utah reports the amount its petroleum storage tank fund paid on each release, and I use it as reported. [Appendix A](#) documents each state’s source field.

From these claims I fit a model mapping a facility’s traits to its expected cleanup cost. Every facility then carries a fitted severity whether or not it ever filed a claim. Two features of cleanup costs determine the estimator. Many releases cost the fund nothing, because the cleanup falls below the deductible, and the claims that do pay have a heavy right tail. Poisson pseudo-maximum-likelihood handles both ([Santos Silva and Tenreiro 2006](#)). A regression of log cost cannot use the zeros, since the log of zero is undefined, and discarding them would select on the outcome. Poisson pseudo-maximum-likelihood retains them. It also estimates the dollar mean $E[C \mid x]$ directly rather than the mean of log cost, avoiding a retransformation that is biased when the error variance moves with the covariates, and it stays consistent under heteroskedasticity.

The features that carry the fit govern how much contaminated material a release puts in the ground and how hard it is to reach. An older tank has been leaking longer before anyone notices, a single wall puts fuel directly against soil, and a larger tank holds more of it. Their interactions matter for the same reason, because a large old single-walled tank is not the sum of three separate penalties. The model is

$$E[C_i \mid x_i] = \exp(\beta_0 + g(x_i) + \alpha_{s(i)}),$$

where C_i is the fund-payable cost of claim i in 2023 dollars, x_i holds the facility’s traits and every two-way interaction among tank age band, single-walled status, tank count, capacity, and the two fuel indicators, $g(\cdot)$ is the function of them an elastic net selects, and $\alpha_{s(i)}$ is a state effect.

Its construction mirrors the hazard model at every step, with a log link and a dollar outcome in place of a logit and an indicator. It uses the same trait set, the same elastic-net selection tuned by cross-validation, and the same out-of-sample scoring. Folds are drawn by facility and stratified by state, and no facility spans training and holdout. Selection does more work here than in the hazard model, because the trait set is rich and the claims sample is three orders of magnitude smaller. Cross-validation again lands at the lasso corner, $\alpha = 1$, and

retains 62 features. A tank inherits its facility’s composition features alongside its own age and wall.

Pennsylvania’s fund supplies an out-of-model check. Its actuarial review implies about \$162,000 per claim and a yearly claim rate near 2.5 to 3 percent (AON, 2022 USTIF Performance Review). The model here fits a median covered severity of \$158,600 and a mean hazard of 2.3 percent for the same state. Both rest on the same claims. The match therefore confirms that the pipeline reproduces the fund’s numbers rather than validating them independently.

Risk rating prices tanks while the fair premium is measured at the facility, and the comparison with filed premiums therefore needs a premium for each tank. I split a facility’s premium across its tanks in proportion to each tank’s share of the facility’s expected loss,

$$P_{ij}^{\text{fair}} = P_i^{\text{fair}} w_{ij}, \quad w_{ij} = \frac{r_{ij} \hat{S}_{ij}}{\sum_k r_{ik} \hat{S}_{ik}}.$$

The relativity r_{ij} is the national age-by-wall hazard entering only as a within-facility ratio, and $\hat{S}_{ij}$ is the tank’s fitted cleanup cost. The split conserves the facility premium by construction, $\sum_j P_{ij}^{\text{fair}} = P_i^{\text{fair}}$, and therefore redistributes risk across a facility’s tanks without re-leveling it. This is ordinary ratemaking practice. A rate is equitable when “differences in rates reflect material differences in expected cost” (Actuarial Standards Board, ASOP No. 12, §3.2.1), and the fair premium is that benchmark.

Both first-stage models carry estimation error, and I carry it into the premium. The two are fit on independent samples, the hazard on the facility-year panel and the cost model on the claims, and I treat their errors as independent, combining them by the delta method, $\text{Var}(h\hat{S}) = \hat{S}^2 \text{Var}(h) + h^2 \text{Var}(\hat{S})$. The hazard variance is the analytic cluster-robust interval the risk model produces at the age-by-wall cell level, and each facility inherits its cell’s hazard uncertainty. The cost variance is a facility-clustered bootstrap. The reported band is $h\hat{S} \pm 1.96 \text{SE}$, floored at zero.


C.2 Filed premiums against the fair premium

The comparison that adds the filed minimum policy premium to the fair premium is [Figure 3](#) in [Section 3.2](#).

[Figure 29](#) shows the same comparison without the minimum-premium adjustment. On the raw comparison the charged premium runs 2.48 times the fair premium pooled, and most of that gap is the minimum policy premium, the floor below which no carrier will write. Weighted across carriers and deflated to 2023 dollars, that minimum is \$563 per policy-year, or \$219 per tank-year split across a facility’s tanks, and it binds for 15.0 percent of Texas facilities. Netting the floor out as the fixed cost of writing a policy puts the implied ratio at 1.42 times the fair premium at the tank and 1.44 at the facility, ranging from 0.67 to 1.55 across age bins. Which of the two comparisons is the right one depends on whether the floor is read as an expense load or as a price, and the filings state it as a floor rather than as a stated fee. Treating it as an additive fixed cost is therefore an interpretation. The

comparison also holds the contract fixed, at coverage above the standard base deductible up to the federal minimum one-million-dollar per-occurrence limit, because facilities' actual deductible and coverage-limit choices are unobserved.

Figure 29: Filed Texas premiums against the fair premium, without the minimum-premium adjustment




Notes: Mean charged premium (blue line) and mean fair premium (black line) per tank-year, by three-year bin of tank age, for Texas tanks in 2015, in 2023 dollars. The figure plots 17,718 tanks. The charged premium applies the filed rating manual of each of five carriers to the insured tanks of the facility. The charged premium is for the standard contract, with coverage between the base deductible and the federal minimum limit of one million dollars per occurrence. The fair premium is the tank's share of its facility's expected cleanup cost, a release rate times a fitted cleanup cost. The fair premium in this figure carries no minimum policy premium, and [Figure 3](#) in [Section 3.2](#) reports the adjusted comparison. Shaded bands are 95 percent intervals. The fair premium band carries the estimation error of the release-rate and cleanup-cost models, and the charged-premium band carries the sampling error of the mean.

The price path over the risk-rated era shows the reform working through composition. In every year from 2006 the matched Texas sample pays far more per tank than the control states pay in fund fees, and that gap is the price of risk the reform introduced ([Figure 30](#)). The Texas per-tank premium falls across the period, from \$1,035 in 2006 to \$573 in 2020, while the control fee drifts down more gently, from \$210 to \$154. The Texas decline is composition rather than a rate cut. As the reform retires the oldest and priciest tanks, the tanks that remain are younger and cheaper to insure, and the average surviving premium therefore falls while its spread across facilities narrows. What a facility pays in total falls


harder still, from \$3,035 to \$1,435, because a facility both sheds tanks, from about 2.93 to 2.51 apiece, and pays less on the tanks it keeps (Figure 31). The level gap is the repricing of risk that Section 4.1 exploits, and the decline over the period is a first look at the closure response that design measures directly.

Figure 30: Per-tank price of coverage, matched Texas sample against the control fee



Notes: Mean annual price of coverage per tank by year, 2006 to 2020, in 2023 dollars. The Texas series (red line) covers the insured tanks of the matched estimation sample, priced under the filed rating manuals of five carriers. Each tank receives the mean premium that its carrier's manual yields for tanks of the same wall construction and five-year age bin in that year. The control series (blue line) is the per-tank fund fee in force, averaged over facilities in the control states of the matched sample. Shaded bands span the 10th to 90th percentiles across tanks for Texas and across facilities for the control states.

Figure 31: Annual premium per facility, Texas against the control fee



Notes: Annual price of coverage per facility by year, 2006 to 2020, in 2023 dollars. The Texas series (red line) covers every Texas facility that one of five carriers prices. Each insured tank receives the mean premium that its carrier’s filed rating manual yields for tanks of the same wall construction and five-year age bin in that year, and the series sums those premiums over the facility’s tanks. The control series (blue line) is the per-tank fund fee in force times the facility’s active tank count, in the control states of the matched sample. Lines are means across facilities, and shaded bands span the 10th to 90th percentiles. The Texas figures are filed premiums and are not amounts recorded on policies.

C.3 What the comparison cannot show

The level of the fair premium rests on pairing a leak frequency from the release database with a claim severity from the fund database, two different event populations. Texas files no claims, and its severity is therefore predicted from the claims states. The severity is a fitted mean over a heavily right-skewed distribution, with a pooled median claim of \$106,000 against a mean of \$348,000. A pure premium is correctly built on the mean, and that is why these premiums look large beside a typical cleanup.

None of this measures welfare. The comparison shows how filed premiums sit against expected cost, and over- or under-charging risk is welfare-ambiguous on its own. Whether a premium that tracks cost raises welfare depends on responses the comparison cannot see. Firms may respond to the premium or ignore it. They may pass it into pump prices, where consumers respond in turn. Nothing in a static comparison settles any of that. [Section 5](#) prices those responses directly.

D Structural Model Details

D.1 Constructing measured revenue

The model needs a measure of what a facility earns, and the data do not provide one directly. I observe no facility’s own gallons sold and no facility’s own prices. What I do observe is every facility’s storage capacity and the public price series that bracket what a gallon of gasoline earns its seller, and I build revenue from those, with one market assumption connecting them: a facility’s throughput moves with its storage capacity. The construction is deliberately coarse, since facilities in the same bin earn different revenues for reasons capacity cannot capture, and I return to what that restriction costs below.

The inputs are public: the retail gasoline price and the state gasoline tax, by state and year, come from the EIA State Energy Data System and the FHWA tax-rate series. The wholesale price refiners charge resellers comes from the EIA refiner sales-for-resale prices, reported by refining region. The federal gasoline tax is \$0.184 a gallon, flat over the sample. Facility size comes from the tank registry, where I sum each facility’s tank capacity and sort facilities into the four bins of [Section 2.3](#). Every input is an administrative price or a registry capacity, which makes the measure available for all facilities in all years, with none of it drawing on a facility’s own accounts. The series cover the 18 study states from 1999 to 2020.

One state’s early-period margin is bounded, not measured. North Carolina sits in the Lower Atlantic, which the Gulf supplies through the Colonial Pipeline, but EIA reports wholesale prices only for the whole East Coast region, which is pulled up by the Northeast. With no sub-region wholesale price published, North Carolina’s margin before 2014 is floored at a small positive value. It is one cell of fifty-four, and the state fixed effect absorbs the rest.

Because most of the retail price of gasoline is the cost of the fuel itself, the retail price is a poor measure of what a station clears. The right measure is the margin: the retail price net of the wholesale price, the state tax, and the federal tax,

$$m_{g,y} = P_{g,y}^{\text{ret}} - P_{r(g),y}^{\text{whl}} - \tau_{g,y}^{\text{st}} - \tau^{\text{fed}},$$

where g indexes the state, y the year, $r(g)$ the state’s refining region, and $\tau^{\text{fed}} = \$0.184$. The retail price and the state tax vary by state. The wholesale price varies by region. The margin $m_{g,y}$ is what a station keeps on a gallon. I hold it in nominal dollars, to match the nominal premium and deductible the facility pays.

Revenue is that margin scaled to the fuel a facility moves, and throughput scales with tank size, and I therefore proxy a facility’s gallons by its capacity bin rather than by volumes I do not observe. Each tank’s capacity is capped at 60,000 gallons before summing, which keeps a few very large facilities from stretching their bin’s mean, and each facility then carries the mean capped capacity of its bin, $\bar{c}_B$. The margin is averaged within each of three rate eras, 1999 to 2013, 2014 to 2018, and 2019 to 2020, written $\bar{m}_{g,e}$, and a facility therefore forecasts a steady local margin rather than a yearly swing. Measured operating revenue is the product, in model units of $\Omega = \$10,000$,

$$R_{B,g,e} = \frac{\bar{c}_B \cdot \bar{m}_{g,e}}{\Omega}.$$

A facility operating in capacity bin B , state g , and era e earns $R_{B,g,e}$. Revenue is thus a small table with one entry per size bin, state, and era, read off the facility’s current bin.

This construction imposes one behavioral assumption: operating revenue is exogenous to the facility’s dynamic problem, a payoff attached to the facility’s state, not a quantity the firm manages. First, revenue is pinned down by three observables, the size bin, the state, and the era, and the only lever a firm has on its revenue is therefore which capacity bin it occupies, and it moves between bins only through the capital choices the model already prices, downsizing and replacement. It has no other lever on revenue: throughput, effort, and retail markup are outside the model. Second, two facilities in the same bin, state, and era earn the identical $R_{B,g,e}$, which assumes away idiosyncratic profitability. There is no unobserved station that quietly earns more in a way that would jointly raise its sales and its survival, which is the endogeneity the Sweeting two-step is built to shut down. Third, because I average the margin within an era, the firm treats the local margin as the permanent flow it keeps, and within-era gasoline-price swings therefore do not drive the timing of exit or replacement. The cost of these restrictions is that any true within-cell revenue heterogeneity is loaded onto the choice shocks and onto ψ rather than modeled. The return is that revenue enters the payoff as a clean, exogenous state, which is exactly what lets ψ be identified from size-by-margin variation rather than from the firm’s own conduct.

D.2 Computing the value function

The linear system $(I - \beta M)V = R$ has one row per state, 298,448 in each of the 104 environments, and the matrix is never formed. The solve iterates the contraction $V \mapsto R + \beta MV$, whose convergence rate is bounded by $\beta = 0.95$, with the matrix-vector product computed in compiled code that applies the action maps, the aging kernel, and the choice weights cell by cell rather than through a stored matrix. Anderson acceleration of depth five replaces each iterate with a least-squares combination of the last five, and it cuts the iteration count substantially relative to plain successive approximation. Each solve runs to an infinity-norm residual below 10^{-9} and halts with an error rather than returning an unconverged value function. The estimator of [Section 5](#) repeats this solve at every parameter evaluation, once per environment, and the counterfactuals of [Section 6](#) repeat it under each policy.

The estimation converges. The nested pseudo-likelihood loop stops at iteration 13, when the parameter vector and the choice probabilities each change by less than 10^{-5} between successive iterations. The choice-probability update is damped at 0.6. The inner maximization at each iteration runs to a tighter tolerance than the outer loop, and the measured step therefore reflects the estimate and not the optimizer’s own precision.

D.3 Model objects and fit

The dynamic model treats a tank’s age as a Markov process. Each year a surviving tank either advances one five-year band or stays, with a probability fixed by the band it currently occupies. [Table 15](#) reports the empirical transition rates behind that process. A tank stays in its own band with probability close to four in five in every band from the first through

the seventh, and it advances with the complementary probability. Band eight holds tanks past thirty-five years. No tank observed there ever leaves, and the kernel makes the band absorbing by construction. This kernel is the facility’s own forecast of its fleet’s aging, estimated once from the registry and imposed on every environment the model solves.


Table 15: Age-band transition kernel

Age band	Stay	Advance
1	0.770	0.230
2	0.769	0.231
3	0.776	0.224
4	0.795	0.205
5	0.812	0.188
6	0.814	0.186
7	0.831	0.169
8	1.000	0.000

Notes: Empirical one-year transition probabilities between tank age bands. The unit of observation is a tank-year. The sample is the 6,949,981 tank-years of the tank-year panel from 1999 to 2020 for tanks that are not closed in the year and have a recorded age. Age is the calendar year less the installation year. Band 1 is ages 0 to 4, each higher band covers the next five years, and band 8 is 35 years and older. Each tank-year is assigned the band of its age and the band of its age plus one year. Stay is the share of a band’s tank-years whose two bands are the same, and Advance is the share whose second band is the next one up. A tank moves at most one band in a year, and Stay and Advance sum to one. Band 8 has no higher band, and its Stay is one by construction.

The portfolio model tracks a facility through a small set of tank cells rather than through each tank’s individual history, and that simplification needs its own justification. [Figure 32](#) reports how often a facility’s active tanks actually differ from one another, by wall type and by installation vintage, at each facility size. Tanks installed at the same station are typically installed together and age together. The median facility carries a zero-year spread in tank age at every size, and most facilities operate tanks of a single wall type. The state space the model solves is therefore not an approximation forced by computation. It matches how stations actually build and replace their tanks.


Figure 32: Within-facility tank homogeneity by size



Notes: Bars show, by facility size, the percent of facility-years with homogeneous tanks. The bars labeled Wall give the percent with all tanks of one wall type, single-walled or double-walled. The bars labeled Age one bin give the percent with the oldest and the youngest tank in the same five-year age band. The bars labeled Both give the percent that meet both conditions. The age bands are ages 0 to 4, 5 to 9, and so on up to 35 years and older. Facility size is the number of tanks at the start of the year: one, two, three, or four or more. The sample is the 2,193,723 facility-years with a recorded tank age in Texas from 2006 to 2020 and in the control states from 1999 to 2020.

The model assigns which tank a facility removes by a fixed rule: the oldest single-walled tank first, then the oldest double-walled tank. [Figure 33](#) and [Table 16](#) ask whether that rule describes what operators actually do, using a conditional logit that compares each removed tank only against the other tanks standing at the same facility in the same year. Wall type and age carry the removal decision. A single-walled tank faces far higher removal odds than a double-walled one, and each additional year of age raises the odds further. Capacity carries no weight at all. Operators remove the tank the post-reform price schedule targets. That is exactly the tank the model's rule names.

Figure 33: Removal odds ratios, within facility



Notes: Odds ratios from the conditional logit of [Table 16](#). Each facility-year with exactly one removal compares the removed tank with the other tanks standing at the facility that year. Each point is the odds ratio of removal for one more year of tank age, for a single-walled tank against any other tank, and for 1,000 more gallons of capacity. Horizontal lines are 95 percent confidence intervals from standard errors clustered by state, and the dashed vertical line at one marks equal odds of removal. The sample is the 27,591 facility-years of the tank-year panel, 1985 to 2020, with two or more tanks and exactly one removal, and the 104,838 tank-years of their tanks.

Table 16: Which tank is removed: within-facility comparison

Closed tank attribute	OR	SE	p
age per year	1.065	0.008	0.000
SW	5.410	0.293	0.000
capacity per 1000 gal	1.000	0.000	0.861

Notes: Conditional logit of which tank a facility removes in a year with exactly one removal. The unit of observation is a tank-year. Each facility-year is one choice set with the removed tank as the outcome and the other tanks standing at the facility as the alternatives. Only differences among the tanks of one facility in one year identify the coefficients, and facility characteristics such as location, demand, owner, and size drop out. The sample is the 27,591 facility-years of the tank-year panel, 1985 to 2020, with two or more tanks and exactly one removal, and the 104,838 tank-years of their tanks. Age is in years, SW is an indicator that the tank is single-walled with every other tank as the comparison, and capacity is in thousands of gallons. OR is the odds ratio, and a value above one means that a tank with more of the variable is more likely to be the one removed. SE is the standard error of the log odds ratio, clustered by state, and p is the two-sided p-value from that standard error.

Table 17: Removal-rule fidelity and the cost of misclassification

Panel A: How often the removal rule names the right tank

Action	Size	N	Match	Wrong wall
Downsize	1	62	1.000	0.000
Downsize	2	1,231	0.784	0.041
Downsize	3	1,706	0.790	0.038
Downsize	4+	5,180	0.709	0.053
Replace	1	580	1.000	0.000
Replace	2	1,195	0.986	0.003
Replace	3	2,737	0.986	0.004
Replace	4+	2,835	0.930	0.014

Panel B: Premium and hazard error on the misclassified events

Action	Size	Misses	Mean premium err	p90 premium err	Mean hazard err pp
Downsize	2	266	58	124	0.210
Downsize	3	359	53	124	0.125
Downsize	4+	1,505	76	159	0.084
Replace	2	17	65	138	0.305
Replace	3	39	55	124	0.168
Replace	4+	198	82	182	0.107

Notes: The structural model assigns which tank leaves a facility by rule: the oldest single-walled tank first, then the oldest double-walled tank. The sample is the 15,526 facility-years from 1999 to 2020 with at least one tank removed and no exit. A downsize installs no tank in the year and a replacement installs at least one. Size is the number of tanks at the start of the year, and 4+ means four or more. Panel A scores the rule against the tanks the facility removed. N is the number of events. Match is the share of events whose removed tanks fall in exactly the wall-and-age-band cells the rule names. Wrong wall is the share of events with a double-walled tank removed and a single-walled tank kept. Panel B covers the events the rule misses, and Misses is their number. Premium errors are absolute differences, in nominal dollars per year, between the premium on the tanks removed and the premium on the tanks the rule removes, with each tank priced on the Texas rate schedule for its wall type and age band in the filing era of the event year. Hazard errors are absolute differences, in percentage points, between the annual leak hazard of the facility's remaining tanks after the rule's removal and after the actual removal, evaluated at their average age band and majority wall type. Mean is the average over a row's events, p90 is the ninetieth percentile, err means error, and pp means percentage points. Source: insurer rate filings for the Texas rate schedule.

A rule that matches the average pattern can still misname the tank in any single event, and the model needs to know how costly that is. [Table 17](#) scores the rule event by event against every observed downsize and replacement. The rule names the correct tank in 84.6 percent

of 15,526 events, the number [Section 5](#) reports for assumption A1. Its errors concentrate in one narrow cell, replacements at four-or-more-tank facilities that shed a single tank, and even there the errors are mostly about which age band the tank falls in rather than its wall type. Panel B prices what those errors cost the model. A misclassified tank carries a small error in its premium and its hazard, and neither error moves the states the model tracks by an amount that would change a facility’s choice.

The model reproduces the aggregate action shares and fits less well within Texas. [Table 18](#) reports the observed and predicted share of facility-years in each action, for the full estimation sample and separately for Texas and the control states. Within Texas the model predicts more downsizing and less exit than the data show. [Section 6](#) discusses that pattern.

Table 18: Model fit: predicted versus observed action shares


	Observed	Predicted
<i>All facility-years</i>		
Maintain	0.978	0.974
Downsize	0.003	0.006
Replace	0.003	0.004
Exit	0.016	0.016
<i>Texas, risk-rated premiums</i>		
Maintain	0.987	0.982
Downsize	0.001	0.011
Replace	0.004	0.005
Exit	0.007	0.001
<i>Control states, flat fees</i>		
Maintain	0.977	0.973
Downsize	0.003	0.005
Replace	0.003	0.004
Exit	0.017	0.018

Notes: Each cell is the share of facility-years in which the group takes the action, and the four shares sum to one within a group up to rounding. Maintain is a year with no removal and no install, downsize a removal of at least one tank with no install, replace a removal with at least one install in the same year, and exit the closure of the last tank. Facility-years that install a tank and remove none are not in the estimation sample. The sample is the estimation sample of the structural model, 1,654,626 facility-years, Texas from 2006 to 2020 and 14 control states from 1999 to 2020. Observed shares are the raw frequencies in that sample. Predicted shares are the model’s choice probabilities at the estimated parameters, averaged over the same facility-years. The shares enter the estimation only through the likelihood and are not moments matched by construction.

D.4 Identification figures


The identification argument of [Section 5](#) rests on two pieces of price variation, and both are first-stage inputs fixed before any dynamic parameter is estimated. [Figure 34](#) plots the premium side, the filed premium a facility faces as its fleet ages, under the Texas risk-based schedules and under the control funds' flat fees. [Figure 35](#) plots the out-of-pocket side, the expected cost a facility carries itself after a release, its annual leak hazard times its state's deductible.

Figure 34: Mean facility premium by fleet age, Texas against the flat-fee control states



Notes: Mean total facility premium per year, in nominal dollars, by the facility's average tank age in five-year bands. Each point is the mean over the facility-years of one state in one age band, with facility sizes pooled, and each line is one state. The heavy dark line is Texas. A Texas facility is priced on the filed rate schedule of the carrier it held that year, or at the priceable-market mean when no schedule can be reconstructed for its carrier ([Section A.6](#)). The other lines are control states, where the premium is the state fund's per-tank fee, averaged over the sample years, times the facility's tank count. The fee does not vary with tank age, and a line at zero is a state with a per-tank fee of zero in the data. The cells shown hold 1,853,068 facility-years with six or fewer tanks at the start of the year, in Texas from 2006 to 2020 and in the control states from 1999 to 2020. Kansas and Maryland and the Texas facility-years that pay no premium to a carrier are excluded. Cells with fewer than thirty facility-years are suppressed. Sources: insurer rate filings for Texas and state fund fee schedules for the control states.

Figure 35: Expected out-of-pocket cost by fleet age, by state



Notes: Mean expected out-of-pocket cost per facility-year, in nominal dollars, by the facility's average tank age in five-year bands. Each point is the mean over the facility-years of one state in one age band, with facility sizes pooled, and each line is one state, labeled at its right end. The cost is the facility's annual leak hazard times its state's deductible, with the hazard evaluated at the facility's average age band and majority wall type. A line's level scales with the deductible, and its slope across age bands comes from the hazard. The heavy teal line is Texas and the orange lines are control states with a state fund. The Texas deductible is the reference deductible in the carriers' filings and not a measured one. Dotted lines at zero are states with no positive deductible on record. The cells shown hold 2,099,962 facility-years with six or fewer tanks at the start of the year, in Texas and the control states from 1999 to 2020. Kansas and Maryland are excluded. Cells with fewer than thirty facility-years are suppressed. Sources: insurer rate filings for the Texas deductible and state fund statutes and rules for the control-state deductibles.

D.5 Welfare accounting

Every term of the welfare identity of [Section 6](#) is a present value per facility in the observed stock. The subscript g names a facility's pricing environment, a fund state or a Texas carrier-year, and the weight $\mu_g(s)$ is the share of the observed facility-years that sit in environment g at state s . The weights sum to one over environments and states. A counterfactual changes the contract an environment carries and leaves the number of facilities that begin in each state unchanged. It changes what they do afterward and how many facilities open.

Each term follows a flow along a facility's own path. For a flow $f(a, s)$ paid when a facility in state s takes action a , let $\bar{f}_g(s) = \sum_a P_g(a | s) f(a, s)$ be its expected value at state s and M_g the transition matrix the choice probabilities induce, with exit absorbing. The present

value of the flow over the observed stock is

$$\text{PV}_g(f) = \sum_s \mu_g(s) [(I - \beta M_g)^{-1} \bar{f}_g](s).$$

A facility that exits stops paying premiums and stops leaking, and the present value counts neither after it leaves. The expected discounted number of leaks is $\text{Leaks}_g = \text{PV}_g(h)$, with $h_a(s)$ the post-action hazard.

Producer surplus is the firm's value in money,

$$PS = \frac{1}{\gamma} \sum_g \sum_s \mu_g(s) V_g(s).$$

Because revenue is exogenous, PS moves across counterfactuals only through how long facilities operate, the fleets they hold, and the costs they bear. A policy raises PS by letting facilities keep their revenue streams at lower cost, and the accounting is silent on demand-side responses. The insurer's position is premium revenue net of claims,

$$\text{Insurer} = \sum_g [\text{PV}_g(\text{prem}) - \text{PV}_g(h S_{aw})].$$

The severity S_{aw} is the fitted mean cleanup cost of a leak, estimated from the claims records of the states that file them and transferred to every facility, Texas included, as the same expectation for the same tank cell. It is capped at the one-million-dollar statutory limit and evaluated at the post-action cell, and a leak at a single-walled fleet therefore books a larger payout than one at a double-walled fleet of the same age. The firm bears the tail beyond the cap,

$$\text{FirmAboveCap} = \bar{S}_{>L} \sum_g \text{Leaks}_g, \quad \bar{S}_{>L} = E[(S - \$1,000,000)^+],$$

and government outlay is nonzero only in the subsidy experiment,

$$\text{Gov} = \frac{1}{\gamma} \sum_g \text{PV}_g(m s_{\text{sub}} c^{\text{inst}}).$$

The last incumbent term is the damage a leak does beyond the cleanup, $\text{External} = E_{\text{ext}} \sum_g \text{Leaks}_g$, charged at the infant-health value of \$25,600 per leak on every contract, the baseline included, and each comparison therefore charges the same damage. The first-best premium prices that term. The other contracts do not, and every leak their facilities retain still imposes it. Under exact internalization the incumbent ledger collapses: premium revenue equals claims plus the above-cap loss plus the external damage, the insurer position equals the last two, and the incumbents' net welfare equals producer surplus less government outlay. That identity holds to machine precision in the first-best run and is a test of the implementation rather than a feature of the results.

The entrant block applies the same flows to the facilities that open. Let D_g be the discounted number of entrants in environment g , the sum over future years of the entry probability times

the pool remaining in that year, and let $\omega_g(s)$ be the distribution of states at which entrants open. An entrant flow is $D_g \sum_s \omega_g(s) [(I - \beta M_g)^{-1} \bar{f}_g](s)$, and the insurer, above-cap, and external terms of the entrant block are built from it exactly as the incumbent terms are. Entrant producer surplus is the value of a new facility and the entry cost is what entrants pay to open,

$$PS^E = \frac{1}{\gamma} \sum_g D_g \bar{V}_g, \quad \text{EntryCost} = \frac{1}{\gamma} \sum_g D_g [ec + (1 - s_{\text{sub}}) c^{\text{inst}} \bar{n}_g].$$

Under the subsidy the government also pays $s_{\text{sub}} c^{\text{inst}} \bar{n}_g$ on each entrant's tanks. [Section D.6](#) gives the closed form for D_g .

The identity rests on two conventions. The money metric divides exactly the payoff-unit terms, producer surplus, the entry cost, and government outlay, by γ and rescales nothing else, because premiums, claims, and the above-cap loss are already dollars. Utility is linear in the coverage-cost dollar, and compensating and equivalent variation coincide. Premiums are a transfer. The firm pays them inside V , the insurer receives them as premium revenue, and only real resource flows remain in W .

Two caveats bound the severity input. Its level is a fitted mean transferred from other states' claims, and because the same S_{aw} enters the baseline and every counterfactual, level error cancels from the parts of each comparison where behavior does not move and scales the parts where it does. Severity never enters a facility's payoff, only the accounting. The second caveat concerns the youngest cells. Leaks in a tank's first decade are mostly installation failures rather than the corrosion the rest of the age gradient measures, and the fitted severity in those cells is correspondingly high. The model books that severity on new tanks. A replacement therefore looks more expensive to insure against than the aging risk alone implies, and the contracts that induce replacement are scored conservatively.

D.6 Entry: the pool, the entry cost, and welfare by pool size

The entry step of [Section 5](#) rests on a pool of potential entrants that is set by rule rather than observed. This section gives the rule, reports the entry cost and how well it fits, and shows how the welfare ranking of [Section 6](#) moves with the size of the pool.

Risk-rated pricing in the fund states ranks first at every pool size. The entry cost rises with the pool, from 12.71 at the baseline pool to 19.60 with no limit on the pool. One entry cost reproduces the total number of entrants and does not reproduce their distribution across states.

The pool. The number of facilities a state can support in a year is λ times its population. The rate λ is the largest ratio of facilities to residents in the sample, counting the facilities operating at the start of a year and the facilities that open during it. That maximum is one facility for every 556 residents, in Arkansas in 1999. The search runs over the control states from 1999 through 2017 and Texas from 2006 through 2017, leaving out Maine and South Dakota, whose registry counts are not reliable enough to define a maximum. The pool of potential entrants in a state and year is the cap less the facilities operating at the start of

the year. South Dakota’s recorded stock exceeds its cap in every year. It is excluded from the entry-cost fit, and in every contract its facilities open at the rate observed in the data.

Entry and the stock of operating facilities determine each other over time, because each year’s pool is the cap less a stock that past entrants have added to. With entry probability p_g , cap N_g , and discount factor β , the discounted number of entrants in environment g solves that recursion in closed form,

$$D_g = \frac{p_g \beta (N_g / (1 - \beta) - \hat{s}_g)}{1 + p_g \beta \hat{S}_g},$$

where $\hat{s}_g = \sum_s \mu_g(s) [(I - \beta M_g)^{-1} \mathbf{1}](s)$ is the discounted number of years the incumbents remain open and $\hat{S}_g = \sum_s \omega_g(s) [(I - \beta M_g)^{-1} \mathbf{1}](s)$ is the discounted number of years an entrant remains open. The first cohort of entrants opens one year after the observed stock is measured.

The entry cost. Table 19 reports the entry cost at the baseline pool and at pools two and five times as large, and at the limit in which the pool never binds, where the binomial likelihood becomes a Poisson one. A larger pool lowers the entry probability that matches a given count of entrants and therefore raises the estimated cost. The standard errors are clustered by state and take the value of a new facility as given.

Table 19: Entry cost by size of the pool of potential entrants

Pool	Entry cost	Standard error	Dollar equivalent
Baseline pool	12.711	(0.893)	\$230,944
Twice the baseline	13.731	(0.868)	\$249,476
Five times the baseline	14.802	(0.859)	\$268,932
No limit on the pool	19.605	(0.854)	\$356,193

Notes: Each row is a separate maximum-likelihood estimate of the entry cost on 259 state-years in 14 states, in which 12,888 facilities open, with standard errors clustered by state in parentheses. Entrants in a state and year are binomial draws from the pool of potential entrants, with a logit entry probability in the value of a new facility net of the install cost of its tanks and of the entry cost. The last row is the Poisson limit of that likelihood. The entry cost is in model units of 10,000 dollars, and its dollar equivalent divides the estimate by γ , the weight on a dollar of coverage cost, and multiplies by 10,000. The entry cost excludes the install cost of the entrant’s tanks. The standard errors take the value of a new facility, the install cost, and the pool as given.

The fitted model matches the 12,888 entrants in total. It does not match where they open. It places 12,829 of them in Texas, where 2,675 facilities opened, and 59 in the control states, where 10,213 opened. The value of a new facility is higher in Texas than in any control state, and with one entry cost and a unit shock scale that difference in value becomes a large difference in the entry probability.

Welfare by pool size. Table 20 splits the net welfare of each contract in Table 3 into its incumbent and entrant blocks and reports the total at each pool size. The incumbent block does not depend on the pool. Entrants are booked at the expected surplus of the entry decision, as in Section 6.3. Risk-rated pricing in the fund states ranks first at every pool size.

Table 20: Net welfare of incumbents and entrants, and totals by pool size

Scenario	Baseline pool, by block		Total net welfare, by pool size		
	Incumbents	Entrants	Baseline	Twice baseline	Five times baseline
(FB) First best	57,632	227	57,859	57,770	57,768
(CF0) Observed contracts	32,109	-2,571	29,539	28,388	27,648
(CF1) Flat fee in Texas	32,335	-5,149	27,186	23,339	19,767
(CF2) Risk-rated pricing in the fund states	45,059	-2,548	42,511	41,355	40,618
(CF3) Replacement subsidy	31,507	-2,707	28,800	27,541	26,709
(CF4) Age-20 removal mandate	27,760	-3,520	24,240	23,528	23,086

Notes: Present value of net welfare in 2023 dollars per facility in the observed stock. Every row charges the same infant-health damage of 25,600 dollars per leak. Incumbents and Entrants are the two blocks of the welfare identity of Section 6.3 at the baseline pool of potential entrants, with entrants booked at the expected surplus of the entry decision. They sum to the Baseline column, the net welfare Table 3 reports. The remaining columns re-estimate the entry cost with a pool two and five times the baseline and re-solve entry under each contract. The incumbent block is the same in every column.

What the entry step cannot show. The pool is an assumption, and the entry cost and the size of the entrant block move with it. The standard errors do not carry the sampling error of the value of a new facility. The welfare ledger counts no benefit of a new facility to the drivers it serves, and the entrant block is therefore a lower bound on the value of entry. A model in which the entry cost or the value of a new facility varies by state would be needed to match where facilities open.