Cleanup Delay and Second-Best Environmental Policy: Structural Estimates from Texas Oil & Gas Wells

Mark Agerton[1], Lauren Beatty[2], and Diógenes Cruz*[3]

[1]Colorado School of Mines

[2]Environmental Defense Fund

[3]UC Davis

November 23, 2026

Draft. Latest version [here].

Abstract

Decommissioning an oil or gas well, plugging it and restoring the site, earns its operator nothing, making delay privately rational. A fee on each damaging well-year reaches the first best without knowing plugging costs, but it must charge idle and producing years separately unless only idle wells do damage or every unplugged well does equal damage. The instruments regulators use instead, such as idling limits, bonds and plugging subsidies, cost more than the fee for the same abatement, and how much more depends on how plugging costs are distributed across wells. However, regulators do not observe these costs, and the available estimates come from orphan wells the state chose to plug, which may not represent the larger population. We recover that cost from the timing of decisions to produce, idle and plug, using a dynamic discrete choice model that accounts for well-level heterogeneity, estimated on low-producing wells in Texas. Our estimates suggest that operators who plug a well pay about $84k in oil and $60k in gas. When only idle wells do damage, a fee on each idle year is the cheapest way to reduce idle well-years. An idling limit of 2 years has 1.29 times the deadweight loss of the idle-year fee that reduces idle well-years by as much in oil (1.35 in gas). A plugging subsidy that abates as much has 2.75 times the fee's deadweight loss (2.97 in gas). For the limit to raise welfare, an idle well-year must do at least $9.2k of damage in oil, 37 times its methane damage.

*Corresponding author: dcruzf@ucdavis.edu, diogenescruz.com.

1 Introduction

Environmental remediation yields no direct profit for producers, so delay is privately rational and the environmental costs of inaction fall on others. The result is a growing stock of undecommissioned facilities, such as open mines, abandoned platforms and unplugged oil and gas wells, the last of which is the focus of this paper. When an oil or gas well stops producing, its operator can keep it idle, return it to production later, or plug it and restore the site, which is permanent. In the U.S., state and federal regulations typically require an operator to plug a well, or obtain an extension, once it has been inactive for a year; operators nonetheless delay plugging, leaving hundreds of thousands of wells idle indefinitely. In Texas alone, the Railroad Commission (RRC) documents more than 155,000 wells as "inactive" in December 2025 (RRC, 2026c).[1] The U.S. has drilled over 2 million wells since the 1950s (EIA, 2025), and about 4 million in its history (Cleveland, 2023). Unplugged wells emit methane (Kang et al., 2014; Boothroyd et al., 2016), threaten groundwater (Ho et al., 2018) and pose a hazard to public safety (Raimi et al., 2021).

Since every well must eventually be plugged, the question for policy is not whether wells are plugged but when, and at what cost to the operators. The current mechanism in place requires operators to post bonds guaranteeing proper well closure. A bond is money the operator deposits in an account, or a letter of credit from a financial institution, that the operator gets back if it plugs the well and forfeits otherwise. However, bonds seldom cover the cost of plugging an abandoned well (Ho et al., 2018). In Texas, the bonds collected from operators who abandoned wells in fiscal year 2015 covered only 15.9% of what the RRC spent per well to plug them (SAC, 2016).

We ask two questions. First, which instrument brings plugging forward at the lowest deadweight loss? Second, how large must the external cost of an idle well-year be for each instrument, at each level, to raise welfare? We ask the second question because the external costs of an idle well are only partly measured. Methane emissions are measured, but other externalities, such as the risk that the well becomes an orphan, groundwater contamination and the loss of property value, are not. Therefore, for each instrument we compute the break-even damage: the external cost per idle well-year above which the instrument raises welfare. It requires no damage estimate, and we compare it with the methane damage alone.

The option to idle is what makes the problem hard. Operators may idle a well to wait for higher prices or to carry out maintenance or recompletion work, which is meant to be temporary. However, operators often idle wells for long periods only to delay plugging and abandonment (P&A) (Muehlenbachs, 2015). In Texas, a well inactive for twelve months must be plugged or hold a plugging extension. The extension is renewed each year, and most operators obtain it through low-cost options, such as plugging a tenth of their inactive

[1] Title 16 of the Texas Administrative Code (TAC), §3.15, as in force in December 2025, defines an inactive well as an unplugged well that has not had any permitted activity for 12 months or more.

wells each year or posting a bond (Section 2), so in practice the rule does not stop long delays. Further, though Texas requires a bond for every well, operators may post a single "blanket bond" that covers many wells, often state-wide (NCSL, 2022). The Sunset Advisory Commission (2016) reported that in fiscal year 2015 the RRC collected $2,707 per well in bonds from abandoned wells, whereas it spent $17,012 per well to plug them, a figure that excludes site remediation.

A policy that brings plugging forward has to be designed around what plugging costs. The RRC publishes a cost calculation for plugging an inactive well, based on what it paid to plug abandoned wells over the previous fiscal year, which ranges from $5.96 to $46.48 per foot of depth (RRC, 2026a). For a well of the mean depth of our oil panel, 4,680 feet, that is $28k to $218k. However, these costs reflect two selection conditions. First, these wells' operators chose to default on their obligations, whether through bankruptcy or bond forfeiture. Second, the RRC chooses which orphan wells to plug, prioritizing those that pose a high risk to the environment.[2] The cost estimates of Raimi et al. (2021) share this limitation: they come from orphan wells plugged by contractors for state regulators.

Hence, the costs published by the RRC are not representative of the broader inactive well population. Wells also differ from each other in ways neither the RRC nor we observe: some are cheap to return to production and others are not, and some have failed and cannot produce again without a costly repair. We address both problems by estimating a structural model of production, idling and plugging decisions for oil and gas wells. Our dynamic discrete choice model captures the option value of leaving a well idle: idling conserves production capacity while prices are low and lets the operator restart when they rise. A static model of the decision to plug cannot capture this option value, and a regression of plugging on prices describes past behavior but cannot predict behavior under a policy that changes the value of waiting. The model also allows for well-level heterogeneity in two features the operator knows and we do not: a permanent type that sets how costly the well is to return to production, and to a lesser degree to plug, and a hidden failure state. With these cost estimates, we compare the welfare effects of price instruments (a plugging subsidy and an idling tax) with those of quantity instruments (a capacity-based plugging mandate and a time-based idling limit). We also treat a bonding requirement as a yearly cost on every unplugged well.

Our estimates suggest that operators who plug a low-producing well pay about $84k in oil and $60k in gas, net of royalties and severance taxes. This cost includes everything the operator bears when it plugs, site cleanup included. Raimi et al. (2021) report median costs of $20k for decommissioning and $76k with land reclamation, for orphan wells the state plugged.

Our headline results take the damage of a producing well-year as zero, because the question of this paper

[2]The RRC has a well plugging priority system in place, based on well completion, wellbore conditions, well location with respect to sensitive areas as well as unique environmental, safety or economic concerns. The weights are in `https://www.rrc.texas.gov/media/3hjhjroj/well-plugging-prioritization.pdf`.

is the idle well that may become an orphan. Damage then depends only on whether a well is idle, and a fee on each idle well-year is the cheapest way to achieve any given reduction in idle well-years; set equal to the damage, it is the first-best policy. We measure the cost of a policy by its deadweight loss: the fall in the sum of the operators' value and the government's revenue, before any reduction in damage is counted. A tax or a subsidy moves money between operators and the government without destroying it, so the deadweight loss is the operator value lost because the policy changes when wells are idled, restarted and plugged. Take the limit of 2 years on idling, which is close to the current Texas twelve-month rule enforced without its extensions. The limit removes 94% of the discounted idle well-years in oil, and the fee that achieves the same reduction has a deadweight loss of $1.50 billion. The limit itself has 1.29 times that loss, and a plugging subsidy that achieves the same reduction has 2.75 times that loss (in gas, 1.35 and 2.97 times). A capacity mandate, which requires every well, producing or idle, to be plugged once its production capacity falls to 3.3 barrels a month or less (196 MMBtu a month in gas), costs 9.2 times the fee that achieves the same, smaller reduction in oil, and 9.8 times in gas.

For the limit to raise welfare, the external cost of an idle well-year must be at least $9.2k in oil and $14.3k in gas, whatever the source of the damage; that is 37 and 57 times the methane damage of an idle well-year, $0.25k, the one damage for which the literature gives a value. At that methane damage, the first-best fee is $0.25k per idle year, reduces discounted idle well-years by 3.0% in oil and 1.9% in gas, and raises welfare, while the limit changes welfare by −$1.9 billion in oil and −$2.0 billion in gas. A small enough tax or subsidy passes this test at any positive damage, because its deadweight loss grows with the square of its level; the quantity instruments have no small dose. A bond, which we treat as a yearly compliance cost on every unplugged well, behaves as a plugging subsidy, and its break-even damage is $19.3k per idle well-year in oil. If producing and idle years do equal damage, the plugging subsidy is cheaper than the idling tax, because the tax returns idle wells to production. A planner who uses a model without the heterogeneity sets the tax or subsidy at 2.6 to 6.7 times the level the estimated model calls for and multiplies the deadweight loss by 1.7 to 31, in part because it overshoots its target.

We contribute in three ways: well-level plugging costs recovered from operators' own decisions, with unobserved heterogeneity; a comparison of price and quantity instruments when idling has option value and well types are unobserved; and break-even damages that need no damage estimate.

The question of whether to use a price or a quantity instrument goes back to Weitzman (1974), who shows that uncertainty in costs breaks the symmetry between them. Stavins (1996) extends the analysis to correlated cost and benefit uncertainty, while Newell and Pizer (2003) consider stock externalities rather than flows. The distinction matters here: idle wells are a stock, and their methane leaks are a flow externality. Our setting differs from Weitzman (1974) in where the asymmetry between prices and quantities comes

from. The regulator is not uncertain about costs in the aggregate; she knows the distribution of wells but cannot condition an instrument on a well's type, so the comparison is between a uniform price and a uniform standard. Figure 1 builds this intuition with a simple model of when to plug. When the damage of a well-year is constant, a fee on the damaging well-year is first-best whatever the heterogeneity, because the wells with the least to give up plug first. A uniform mandate or time limit cannot select wells this way and loses more the more dispersed their costs are. Since the heterogeneity is in costs, we expect a price instrument to do better. A fee on idle years loses its advantage only when it misses a damaging margin, as when producing wells also leak. We nonetheless compare quantity instruments, because they are the ones regulators use: Texas regulates inactive wells by time idle, and in 2025 its legislature tightened those rules: from September 2027, a well inactive for more than 15 years and completed more than 25 years earlier cannot obtain a plugging extension, except under a compliance plan, a finding of financial hardship or a full-cost bond (Texas Legislature, 2025).

Our paper also relates to the literature on cost estimation in natural resource industries, which is scattered and seldom addresses the end of an oil or gas well's life. Most of it estimates costs to assess the viability of projects, as Shafiee and Topal (2012) and Guo et al. (2021) do for mining; Rice and Smith (1977) estimate industry-wide parameters of oil and gas supply, including drilling costs, in a static and deterministic model without decommissioning. We instead recover the end-of-life costs these models leave out, at the well level, from the operators' own decision timing.

On the theory side, the optimal timing of environmental cleanup has a literature of its own. Caputo and Wilen (1995) and Lappi (2018) study the optimal rate of cleanup of continuously accumulated waste, and Yang and Davis (2018) and Lappi (2020) extend it to extraction followed by remediation. Yang and Davis (2018) find that a Pigouvian stock tax implements the optimal remediation while a flow tax and a reclamation bond do not, and Lappi (2020) shows that a mandatory shut-down date can also implement the optimum. These results assume deterministic production capacity and fixed prices. Our counterfactuals compare the same kinds of instruments, a fee, a bond and a deadline, when prices and capacity are stochastic, so idling carries an option value, and when wells differ in costs the regulator does not observe.

Our dynamic discrete choice model follows Rust (1987). Closer to our setting, Muehlenbachs (2015) estimates a dynamic discrete choice model of production, idling and decommissioning of oil and gas wells in Alberta, Canada, and shows that idling is often used to delay or avoid environmental remediation. We construct our model in a similar manner and additionally show which features of the data determine each estimated parameter, by inverting the conditional choice probabilities type by type (Hotz and Miller, 1993), using plugging as a terminal action in the manner of Arcidiacono and Miller (2011). While Muehlenbachs (2015) conducts counterfactuals on alternative prices and technological change, our counterfactuals are policy-

oriented, comparing price and quantity instruments.

Unlike these models, ours allows for unobserved heterogeneity. Among idle wells, those that return to production do so early, and the wells that remain idle for many years are mostly those that are expensive to restart or can no longer produce, a form of self-selection (Heckman and Singer, 1984). We allow for two persistent sources of heterogeneity. The first is a permanent reactivation-cost type, which we integrate out of the likelihood by quadrature, as in mixed logit (Train, 2009) and the discrete types of Keane and Wolpin (1997). The second is a hidden failure state that the operator knows and we do not; we compute its probability for each well and year with a forward filter over the well's record, and its identification rests on Hu and Shum (2012). Arcidiacono and Miller (2011) show how to estimate such models from conditional choice probabilities. Without the heterogeneity, the realized plugging cost in oil is \$40k instead of \$84k, and a planner who uses that model overshoots its target (Section 5.6).

The paper is structured as follows. Section 2 describes the policy setting, with a simple model of why a limit on time idle costs more than a fee on idle years, and our data, the sample of low-producing wells on which we estimate, and the descriptive evidence. Section 3 presents the structural model, its unobserved heterogeneity and how we estimate it. Section 4 gives the estimates and the fit of the model. Section 5 compares the policy instruments, and Section 6 concludes.

2 Policy context and data

2.1 The life of a well and its regulation

The oil and gas industry accounted for about 0.8% of U.S. GDP in 2025 (BEA, 2026). Throughout its history, the U.S. has drilled over 4 million oil and gas wells across the country, and many more test, injection and service wells (Cleveland, 2023). The number of producing wells peaked in 2014 with over one million wells, but has since declined to just below 920,000 wells in 2023 (EIA, 2024).

The life cycle of a well may be roughly divided into five stages:

- **Exploration:** The process of searching for oil and gas reserves, which may involve geological surveys, seismic studies, and exploratory drilling.
- **Drilling:** The actual drilling of the well, which consists of drilling holes, casing the well to protect aquifers, and the installation of safety equipment. This stage may include the prior preparation of the surface, including clearing the land and building access roads.
- **Completion:** This is the process of preparing the well for production, which may include perforating the casing, fracturing the rock, and installing production equipment.

- **Production:** Production refers to the actual extraction of oil and gas from the well, which is then transported to refineries or processing plants through pipelines and trucks.

- **Plugging and Abandonment:** This is the final stage of the well life cycle, when the well is no longer economically viable. This stage includes the removal of all equipment and the sealing of the well with cement and other materials to prevent leaks and groundwater contamination. Further, decommissioning of the site may include the removal of all surface equipment, restoration of the land, and monitoring for any potential environmental impacts. P&A is permanent.

An operator may stop producing from a well before the end of its life. This may be a temporary measure to conserve the oil and gas while waiting for higher prices or to carry out recompletion works. Often, however, operators stop production to delay plugging rather than to keep an option to produce later, creating a mounting stock of idle wells, some of which are at risk of becoming orphans.[3]

We study Texas because oil and gas matter to its economy and the nation's, because it has many idle and orphan wells, and because the data are available. According to the Energy Information Administration (EIA), Texas is the largest producer of oil and gas in the U.S., with 42% of U.S. crude oil production and 28% of U.S. marketed natural gas production in 2023 (U.S. Energy Information Administration, 2025a,b), and 23% of U.S. gas wells that year (U.S. Energy Information Administration, 2025d). The Railroad Commission counts 434,091 wells in Texas in December 2025, of which 155,562 are inactive. Of its 157,151 active oil wells, 99,378 produce less than 10 barrels of oil per day, and of its 84,134 active gas wells, 62,480 produce less than 250 MCF per day (RRC, 2026c). Texas is therefore a suitable setting in which to estimate structurally what it costs operators to produce, idle and plug wells.

A Texas well becomes inactive after twelve months with no reported production, injection or other permitted activity, and its operator must then plug it or hold a plugging extension.[4] The extension is checked when the operator renews its organization report (form P-5) each year, so the first compliance cost falls between twelve and twenty-four months of inactivity. In our yearly data this is a limit of two idle years, $\bar{\kappa} = 2$ in Section 5. Separately, a tax incentive in force through February 2010 exempted from severance tax, for ten years, a well that had been certified as inactive for two years and returned to production.[5]

To encourage compliance with plugging requirements, Texas has two bonding requirements in place:

- A baseline bond set at the individual well level for $2.00 per foot of depth or a blanket bond for:
  - $25,000 for 10 or fewer wells;

[3] An orphan well is an abandoned well that has no responsible operator.
[4] 16 TAC §3.15(a) and §3.14(b)(2); Natural Resources Code §89.011; as in force in December 2025.
[5] Tax Code §202.056 (oil) and §201.053 (gas); 16 TAC §3.83; as in force in 2010.

  - $50,000 for up to 99 wells; or
  - $250,000 for 100 or more wells

- For inactive wells, these must have an inactive well plugging extension, which may be a blanket extension or individual extension, whose compliance options are shown in Table 1, along with the number of organizations and wells selecting into each option.

A blanket bond covers all of an operator's wells for a fixed amount that depends only on the number of wells. Plugging one well therefore does not lower what the operator posts, except when it moves the operator to a lower tier, and the operator recovers nothing until all its wells are plugged. The marginal benefit of plugging one well, through the bond, is therefore zero. We observe selection into forms of compliance, with most operators choosing a blanket extension of either P&A 10% of the wells, or the bond of up to $2 million. This selection into different W-3X certification options reflects operators' private information about their well-specific plugging costs, operational capacity, and financial constraints. We do not model selection into compliance options; the distribution of well types in the model absorbs it. Our structural model estimates costs from a different margin: the timing of decisions to produce, idle or plug across the state space. This in turn reflects average costs across compliance types, which are relevant for the policy counterfactuals we conduct.

Table 1: W-3X certification options in Texas

Option	Organizations	Wells
Blanket extension		
A: plug 10% of inactive wells each year (wells of record at the last P-5 renewal)	559	18,788
B: publicly traded entity, RRC as secured creditor for asset retirement funds	0	0
C: plugging bond, the lesser of the cost to plug all inactive wells and $2 million	163	33,657
Individual extension		
D: abeyance of plugging report	387	3,275
E: Commission-approved enhanced oil recovery	54	2,847
F: fluid level test (H-15)	64	182
G: supplemental financial assurance equal to the plugging cost	22	33

Wells counted by unique API number.

The extension requirements in their current form come from HB 2259 (2009), which the Railroad Commission implemented in September 2010. Since then each extension requires one of the financial-assurance options of Table 1. Further, to obtain an extension for a land well whose surface it does not own, the operator must have purged the well's equipment of production fluids once the well has been inactive for five years, and removed its surface equipment once it has been inactive for ten.[6] In 2025 the legislature added that, from

[6] 16 TAC §3.15(f)(2)(A), as amended effective January 1, 2017; as in force in December 2025.

September 2027, a well inactive for more than 15 years and completed more than 25 years earlier cannot obtain an extension, except under a compliance plan to plug or reactivate it by 2042, a finding of financial hardship, or a bond for its full plugging cost (Texas Legislature, 2025). HB 2259 falls in the middle of our sample and is followed by a wave of plugs in 2011–12. Our model has one regime over the whole sample and does not reproduce the wave (Section 4). The RRC's cost calculation (RRC, 2026a), which comes from the orphan wells it chose to plug, also sets the amount of financial assurance that operators of inactive wells must post.[7]

A natural response to the bonding shortfall is to raise bond levels. However, the literature on environmental bonds identifies several obstacles. First, bonds impose transaction costs and liquidity constraints, particularly for smaller operators (Gerard, 2000; Boyd, 2001), and widespread use of blanket bonds undermines their efficacy (Ho et al., 2018; Boomhower, 2019). Further, while some authors find that full-cost bonding is optimal (Costanza and Perrings, 1990; Kahn et al., 2001), others argue that optimal bonding may be below the full cost of remediation due to reputation effects (Gerard, 2000), marginal-cost considerations (Davis, 2015) or government revenue concerns (Yang and Davis, 2021). Empirically, Boomhower (2019) finds that increased bonding in Texas reduced orphan wells by 70% at the cost of driving out small operators. Whether a bond, or any other instrument, is worth its cost depends on what plugging costs operators and on how they respond, which our structural model measures. We study a bonding requirement in Section 5 as a yearly cost on every unplugged well. To the operator, that cost is the premium paid to the financial institution that issues a letter of credit, or the interest forgone on money deposited with the regulator. No one in the market we study receives it, so we count all of it as a cost.

The Texas rules act on time idle, a quantity, whereas a fee on idle years would act on its price. Figure 1 shows why the difference matters, in a simple case with idle wells that will never produce again. Keeping well i unplugged costs its operator θ_0 a year, and plugging it costs $\psi_{es,i}$ once. Spread over every year after the plug, the plugging cost is $(1-\beta)\psi_{es,i}$ a year, so plugging costs the operator $(1-\beta)\psi_{es,i} - \theta_0$ a year net of the idle cost it ends, and an operator for whom this is positive never plugs. If an idle well-year does damage d_s, plugging well i in year T instead of never changes welfare by

$$\frac{\beta^T}{1-\beta}\Big[d_s - \big((1-\beta)\psi_{es,i} - \theta_0\big)\Big], \tag{1}$$

the discounted idle years the plug removes times the damage of a year less its net plugging cost. The first best plugs at once every well whose net plugging cost is below d_s and leaves the others. In the figure, wells are sorted by their net plugging cost along the horizontal axis and their idle years run up the vertical axis,

[7] 16 TAC §3.15, as in force in December 2025.

each year weighted by its discount factor, so equal areas are equal reductions in discounted idle well-years. A fee τ_s on each idle year makes every well whose net plugging cost is below τ_s plug at once, and removes a vertical strip of the cheapest wells to plug. Setting $\tau_s = d_s$ it is the first best. An idling limit of $\bar{\kappa}$ years makes every well plug after $\bar{\kappa}$ idle years, and removes every well after the deadline, regardless of its plugging cost. When both remove the same area, the fee removes the idle years of the wells that are cheapest to plug and the limit removes the same share of every well's years, so the limit costs operators more. In the figure, with net plugging costs spread uniformly, it costs 3.3 times as much. The gap grows with how much net plugging costs differ across wells, and it vanishes when all wells are alike.

This simplified example leaves out what makes idling worth something: the option to return to production. The structural model in Section 3 introduces this option, along with allowing wells to differ in their reactivation and plugging costs.

Figure 1: An idling limit and an idling fee that remove the same discounted idle well-years

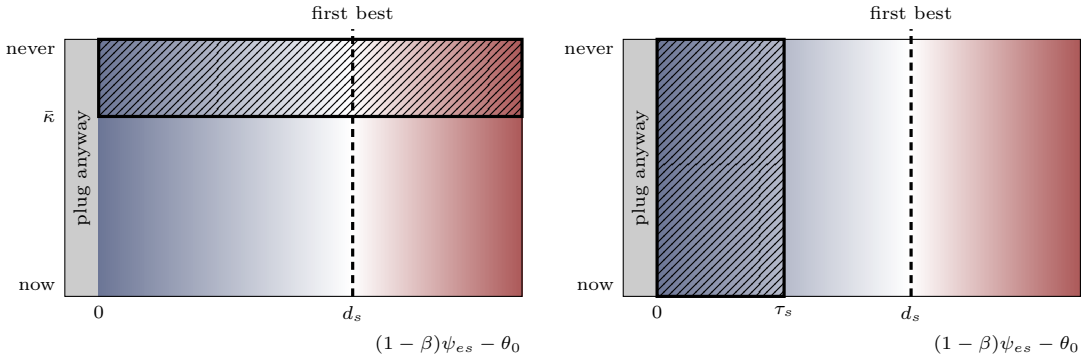

(a) Idling limit $\bar{\kappa}$: every well plugs at $\bar{\kappa}$

(b) Idling fee τ_s: wells below τ_s plug now

Horizontal axis: wells sorted by net plugging cost per year. Vertical axis: idle years, year t at $1-\beta^t$. Hatched: idle years removed. Shading: welfare gain (blue) or loss (red) from removing a year, equation (1). Wells with $(1-\beta)\psi_{es,i} \leq \theta_0$ plug under any policy (grey, not to scale).

2.2 Data

We gather information on wellbores from the RRC: when each was first completed, how deep it is, and whether and when it was plugged (Railroad Commission of Texas, 2025a,b). We take monthly oil and gas production by well from Enverus (Enverus, 2026), along with several well attributes. The RRC records oil production by lease rather than by well, so on a lease with several oil wells the monthly oil volume that Enverus reports for a well is Enverus's allocation of the lease volume. Enverus marks a well as shut in only after three consecutive zero well tests, so an idle well on a lease that keeps producing can carry a small allocated volume for some time.[8]

[8]In Section 4.4, we re-estimate the oil model on the wells that are alone on their lease, whose production is the lease's own and is not allocated. The estimated plugging costs are lower on these wells, and the instruments rank in the same order.

Both sources identify wells by their American Petroleum Institute (API) number. The first ten digits of this number identify a unique well (two digits for the state, three for the county and five for the well), and longer versions of the number add digits that distinguish sidetracked wellbores and completions of the same well. We treat each 10-digit API number as a well, and where a source holds several records for the same 10-digit number, as the Enverus well data can, we combine them. A well is in our data only if it appears in the RRC's wellbore data, which leaves out potential wells that were permitted but never drilled or completed. For each well, the data have two parts: time invariant well-level attributes, and oil and gas production recorded monthly. The attributes are the completion date, the plugging date and the depth from the RRC and, from Enverus, the well type (oil, gas, or both), the drill type, and the true vertical depth, which we use where the RRC gives no depth. We merge the attributes onto the monthly production records by API number.

We keep wells first completed between 1970 and 2009. Records for wells completed before 1970 are often incomplete. Wells completed after 2009 are increasingly horizontal. This restriction focuses our attention on older vertical wells with accurate data. We drop wells with no completion date, and we keep wells whose RRC type is oil, gas, or oil and gas, which removes injection, disposal and water-supply wells and dry holes. Of these, we keep the wells that Enverus types as oil or as gas and drop those it types as both, and we refer to a well's Enverus type as its fuel. We then filter on drill type, keeping the wells Enverus records as vertical and dropping horizontal and directional wells and wells of unknown drill type. We also drop the wells that the RRC flags as plugged but for which it records no plugging date.

The monthly production data run from January 2000 to December 2019, and we classify each well-month as producing, idle or plugged. A month is without production if the well's oil and its gas are both at or below 0.01 (barrels of oil, MCF of gas) or missing, so an oil well that reports only gas in a month is producing that month. A well is idle in every month of a spell of twelve or more consecutive months without production, counted from the spell's first month, and shorter breaks in production count as producing.[9] A well is plugged in the month of its plugging and leaves the panel afterwards. We take the earliest plugging date on record and ignore any production reported after it.

We then collapse the months to a yearly panel of decisions from 2001 to 2018. In the row for year t, the state is the well's status in December of $t-1$ and the action is its status in December of t, or plugged if the well is plugged during t.[10] A well is therefore idle at the start of year t if December of $t-1$ falls inside a spell of twelve or more months without production. The months of 2000 are used only to label spells already

[9] E.g., if a well had positive production in December 2002 and no production between January 2003 and March 2004, it is idle for all 15 of those months. If it only suspended production from January to March 2003, it is producing during those 3 months.

[10] In a well's first year in the panel, the state is its status in its first month in the panel, which is January 2001 for wells completed before 2001.

under way at the start of 2001, and the months of 2019 let us label the December 2018 status of a well whose spell without production begins late in 2018. The 2019 rows do not enter the estimation. Their state, the December 2018 status, gives the stock of wells at the start of 2019. A well's age is the year minus its completion year.

A well's production capacity in a year is its mean monthly production over the months of that year in which it produced, counting only its own fuel: oil for oil wells, and gas for gas wells, converted at 1.037 MMBtu per MCF. While a well is idle, its capacity is carried forward from the last year in which it produced. Wells that enter the panel already idle have no production in the sample, so we seed their capacity with their mean production over their last (up to twelve) producing months between 1980 and 2000, and we measure how long they have been idle from their last producing month before 2001. We drop well-years without a capacity or a price, such as those of a well with no production of its own fuel between 1980 and 2019. We express capacity in thousand barrels (thousand MMBtu) a year, twelve times the monthly mean divided by 1,000, so that capacity times price is gross revenue in thousand dollars a year.

We take monthly spot prices from the EIA, retrieved through the FRED database of the Federal Reserve Bank of St. Louis: the West Texas Intermediate (WTI) price at Cushing for oil, in dollars per barrel, and the Henry Hub price for natural gas, in dollars per MMBtu. We deflate them to 2019 dollars (the 2019 annual average) with the Consumer Price Index for all urban consumers from the Bureau of Labor Statistics, and the reduced-form evidence and the structural model use the annual average of the deflated monthly price. An oil well faces the WTI price and a gas well the Henry Hub price, whatever else the well produces. Section 3 describes the price process of the model.

The panel holds 160,215 oil wells with 2,369,159 well-years and 92,357 gas wells with 1,337,225 well-years over the decision years 2001 to 2018 (Table 2). At the start of 2019, 116,173 oil wells remain, 37,811 of them idle, and 71,475 gas wells, 18,886 of them idle.[11]

Figure 2 shows where the wells in the panel are at the start of 2019. About half of the oil wells (51%) are in the Permian Basin of West Texas, while gas wells are spread across East Texas (25%), the Gulf Coast (25%), the Permian (23%) and the Fort Worth Basin (15%). Figure 3 follows the 133,650 oil wells and 56,536 gas wells that were already in the panel at the start of 2001. By the start of 2019, 31% of these oil wells and 29% of these gas wells have been plugged, and the share still producing has fallen from 77% to 45% in oil and from 83% to 49% in gas. The share idle changes much less, from 23% to 24% in oil and from 17% to 23% in gas. In the panel as a whole, the number of idle wells grows from 30,842 to 37,811 in oil and from 9,571 to 18,886 in gas, and most of that increase comes from wells that entered the panel after 2001 (74% of

[11] These counts cover only the wells of our panel, vertical oil or gas wells completed between 1970 and 2009. They are therefore far smaller than the 155,562 inactive wells the RRC counts in December 2025 (Section 2), which include wells of every type and vintage and follow the RRC's own definition of inactive.

it in oil and 66% in gas).

Figure 2: Producing and idle wells of the panel by county, start of 2019

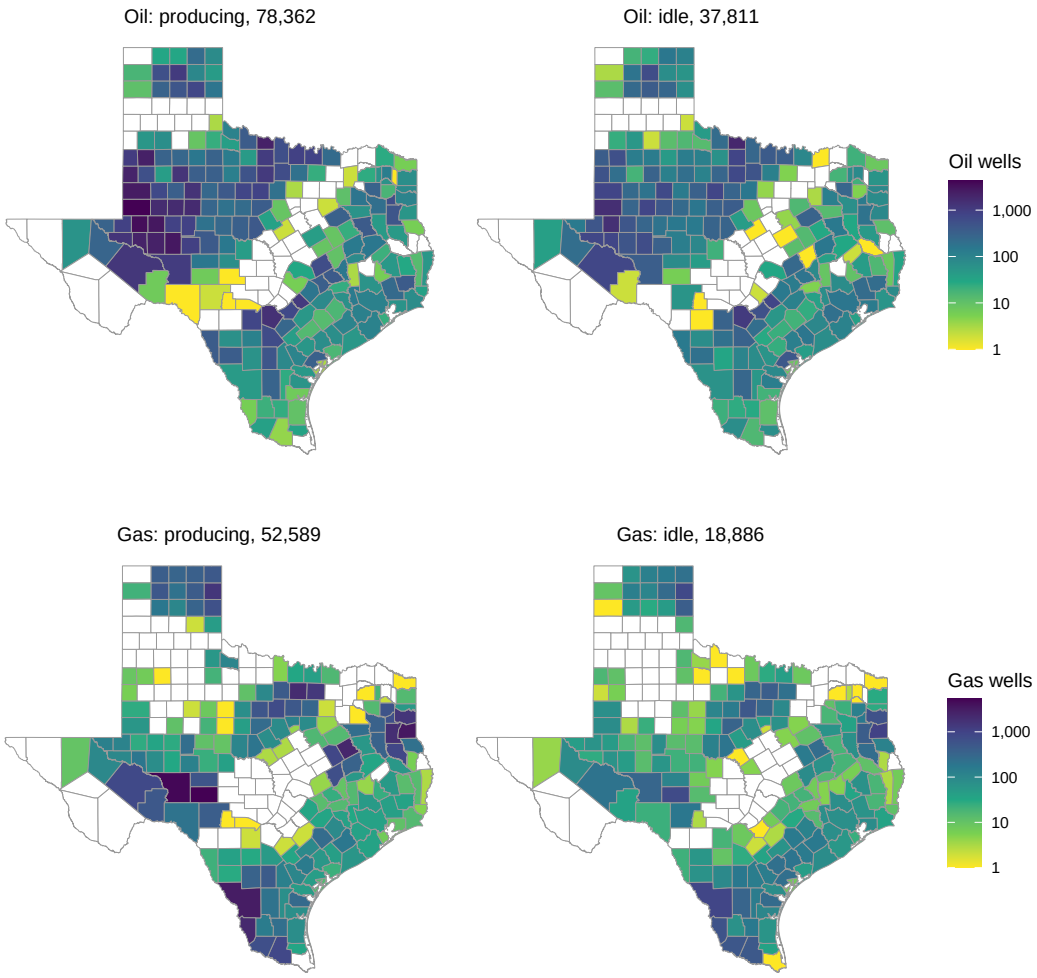

2.3 The estimation sample

We estimate the model on low-producing wells and not on the whole panel. The policy question is when wells near the end of their productive lives get plugged, and wells that still produce a lot are rarely plugged. In the panel, a producing well in the top fifth of the capacity distribution is plugged in 0.2% of its years in oil and 0.2% in gas, compared with 1.8% and 1.7% in the bottom fifth. Well-years at high capacity therefore tell us little about how operators decide when to plug. The restriction is also less narrow than it may seem, because most Texas wells produce little: in December 2025, 99,378 of the 157,151 active oil wells in the state

Figure 3: Status of the wells in the panel at the start of 2001, 2001–2019

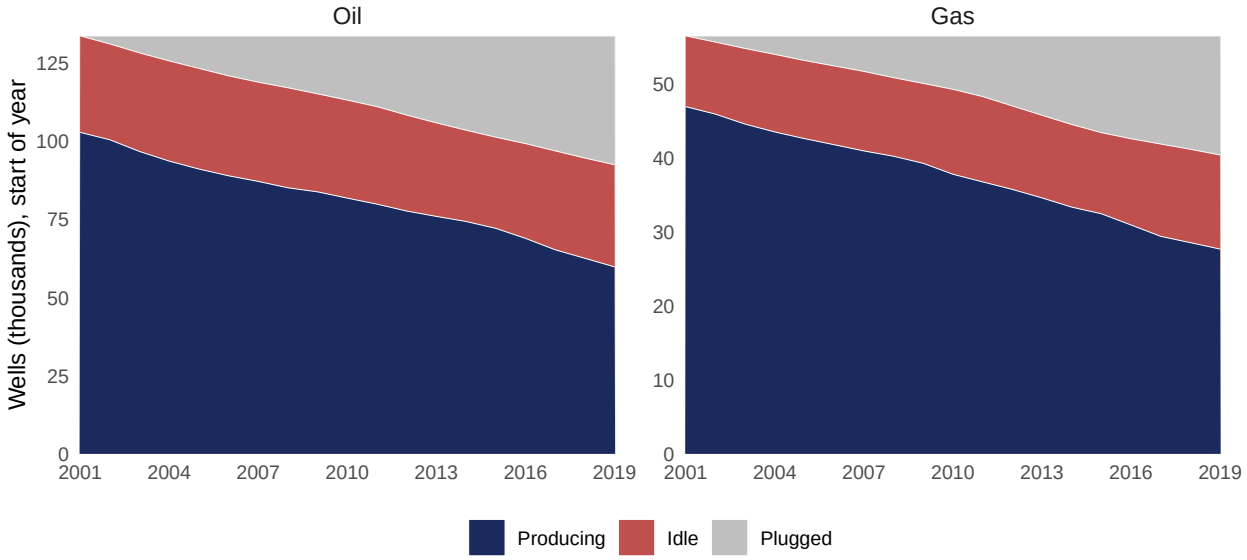

made under ten barrels a day (RRC, 2026c).[12]

For each fuel, we compute the 40th percentile of production capacity across all well-years of the panel, producing and idle. The cutoff is 20.2 barrels a month in oil (0.66 a day) and 758.3 MMBtu a month in gas (24.93 a day). A well enters the estimation sample in its first year at or below this cutoff, which may be its first year in the panel or a later year once its production has declined, and it stays in the sample in every year after that. The sample holds 90,980 oil wells with 1,071,027 well-years and 60,045 gas wells with 598,760 well-years, and it contains 67% of the panel's plug events in oil and 71% in gas. The counterfactuals of Section 5 are for this population only.

We keep every year of a well from its entry on, instead of selecting well-years one by one, because the hidden-state filter of Section 3 reads a well's consecutive rows as consecutive years. Some retained well-years therefore lie above the cutoff, but few do (11.5% in oil and 10.7% in gas): capacity is held fixed while a well is idle, as in the model, and it mostly declines while the well produces.

Table 2 compares the estimation sample with the panel it is drawn from. The wells in the sample are older and shallower than those in the panel, more of them are ever plugged, and more of their well-years are idle (36.2% against 25.0% in oil, and 27.6% against 17.9% in gas). Table 3 gives the yearly transitions between statuses in the estimation sample. A producing oil well is idled in 6.3% of its years and plugged in 1.4%, and an idle oil well returns to production in 4.7% of its years and is plugged in 5.2%. In gas, too, an idle well is plugged more often (6.4%) than it returns to production (4.8%).

[12] Our oil cutoff of 0.66 barrels a day is far below this line. However, a well stays in the sample once it falls to the cutoff, so the sample holds more than half of the panel's oil wells (90,980 of 160,215) and close to two thirds of its gas wells (60,045 of 92,357).

Table 2: The panel and the estimation sample

	Panel	Estimation sample
Oil		
Wells	160,215	90,980
Well-years	2,369,159	1,071,027
Plug events	44,063	29,699
Share of the panel's well-years (%)		45.2
Share of the panel's plug events (%)		67.4
Wells ever plugged (%)	27.5	32.6
Well-years idle (%)	25.0	36.2
Production capacity (BBL/day)	5.0	0.6
	(17.3)	(4.6)
Capacity cutoff (BBL/day)		0.7
Well-years above the cutoff (%)	60.0	11.5
Real price (USD/BBL)	75.10	75.13
	(26.18)	(25.97)
Vertical depth (ft)	4,680	3,695
	(2,920)	(2,700)
Completion year, median	1985	1984
Well age (years)	21.7	25.1
	(10.2)	(8.7)
Gas		
Wells	92,357	60,045
Well-years	1,337,225	598,760
Plug events	20,888	14,756
Share of the panel's well-years (%)		44.8
Share of the panel's plug events (%)		70.6
Wells ever plugged (%)	22.6	24.6
Well-years idle (%)	17.9	27.6
Production capacity (MMBtu/day)	132.5	16.3
	(528.2)	(50.8)
Capacity cutoff (MMBtu/day)		24.9
Well-years above the cutoff (%)	60.0	10.7
Real price (USD/MMBtu)	5.72	5.26
	(2.67)	(2.57)
Vertical depth (ft)	8,296	7,590
	(3,395)	(3,268)
Completion year, median	1998	1995
Well age (years)	15.7	20.2
	(10.9)	(10.4)

Means (standard deviations); years 2001–2018; 2019 dollars.

Table 3: Annual transition probabilities in the estimation sample

	Produce	Idle	Plug	Well-years
Oil				
From producing	0.923	0.063	0.014	683,666
From idle	0.047	0.901	0.052	387,361
Gas				
From producing	0.928	0.062	0.010	433,497
From idle	0.048	0.888	0.064	165,263

2.4 Descriptive evidence

On the estimation sample, we run a multinomial logit of each year's action on the log of the well's production capacity, the log of the year's average real spot price (Cushing West Texas Intermediate for oil and Henry Hub for natural gas), and the log of the well's true vertical depth. Capacity is the well's mean monthly production in the year, over the months it produced, and is carried forward while the well is idle, as in the structural model. We run it separately for producing and idle wells and for each fuel, with the current status as the base category, so the coefficients describe the four switches we observe: producing to idle, producing to plugged, idle to producing and idle to plugged. Table 4 reports the results.

Table 4: Multinomial logit of the yearly choice

	From producing		From idle	
	Idle	Plug	Produce	Plug
Oil				
Log capacity	-0.276***	-0.128***	0.468***	0.102***
	(0.002)	(0.004)	(0.006)	(0.003)
Log price	-0.367***	-0.199***	0.291***	-0.022
	(0.013)	(0.028)	(0.021)	(0.020)
Log vertical depth	0.025***	0.525***	-0.443***	-0.148***
	(0.007)	(0.015)	(0.012)	(0.009)
Constant	-2.373***	-7.973***	1.008***	-1.122***
	(0.072)	(0.172)	(0.133)	(0.110)
Well-years	683,554		386,805	
Gas				
Log capacity	-0.606***	-0.546***	0.295***	0.032***
	(0.007)	(0.010)	(0.011)	(0.006)
Log price	0.022	0.284***	0.173***	0.004
	(0.016)	(0.033)	(0.026)	(0.022)
Log vertical depth	0.622***	0.349***	0.284***	-0.117***
	(0.018)	(0.031)	(0.024)	(0.015)
Constant	-7.748***	-7.621***	-5.800***	-1.633***
	(0.162)	(0.287)	(0.222)	(0.141)
Well-years	433,497		165,263	

Base: stay. SEs clustered by well. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

In Table 4, the coefficient on capacity is negative and statistically significant for the switches from producing to idle and to plugged, for both oil and gas wells, consistent with higher-producing wells being less likely to be idled or plugged. For idle wells, the coefficient on capacity is positive for the switch from idle to producing, in both fuels, which aligns with our expectation that higher potential production from a well makes it more likely to be reactivated. The coefficient on plugging from idle is also positive, though much smaller: among idle wells, those with more capacity are more likely to leave idle status by either route.

The results from the regression also show the expected direction of the coefficients for prices on oil wells. A higher price decreases the probability of idling or plugging a producing oil well and increases the probability of reactivating an idle one; it has no statistically significant effect on plugging from idle. For gas wells, higher prices increase the probability of reactivating, but they also increase the probability of plugging a producing

well, which is hard to rationalize.

Finally, deeper oil and gas wells are associated with a higher probability of being idled or plugged while producing, and are less likely to be plugged once idle. Once we control for price and production, this reflects the higher operating costs (e.g., lifting costs) of deeper wells, which make them less profitable, and hence more likely to be idled or plugged; the lower plugging rate of idle deep wells is consistent with their higher plugging costs. Deep idle oil wells are also less likely to be reactivated, while deep idle gas wells are more likely to be.

We also obtain the CCPs for quantiles of capacity, prices and true vertical depth. Figure 4 shows that the conditional choice probabilities of continuing to produce are increasing in capacity both for oil and gas, while the probability of remaining idle falls, mostly in the top quintile. For producing wells, the probability of switching to idle is decreasing in capacity, while for idle wells, the probability of returning to production is increasing in capacity (only slightly for gas). Probabilities of plugging do not vary much with capacity.

Figure 4: Conditional choice probabilities for each fuel by quantiles of capacity

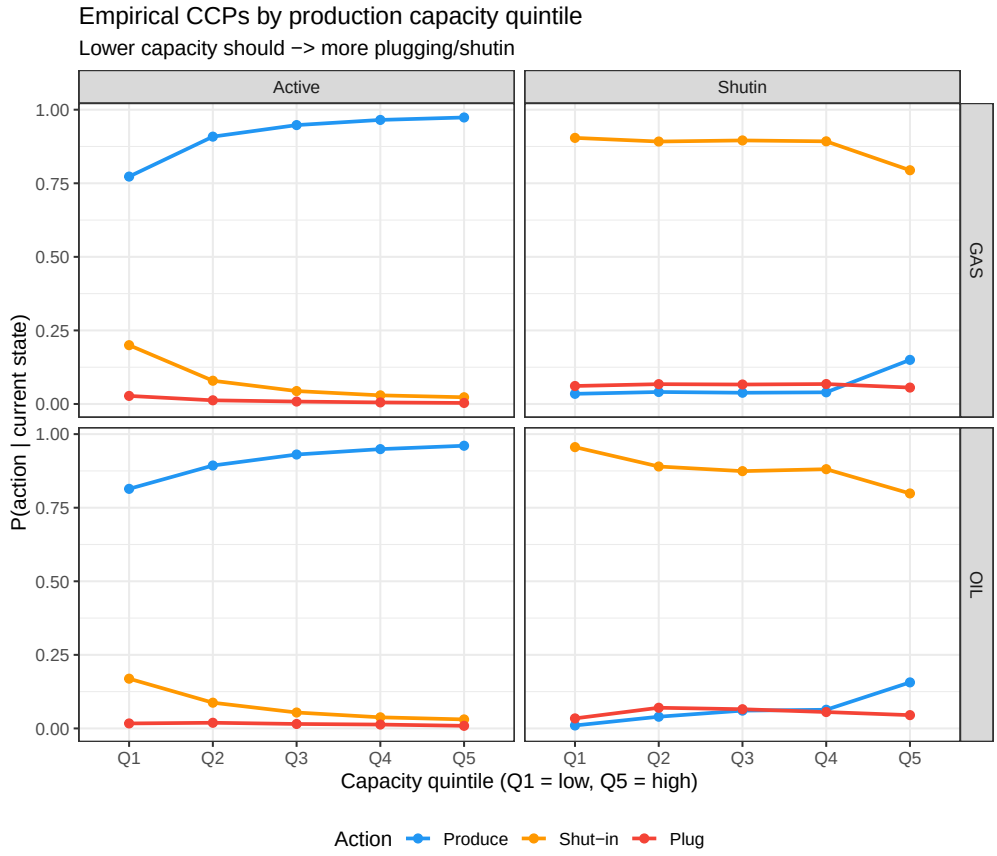

Figures 12 and 13 in the appendix show the CCPs for quantiles of price and depth. While price does bring some variation to the model which can be exploited for identification, CCPs by quantiles of depth show no

significant variation to help identify the model, which is why depth is not a state variable of the structural model.

The longer a well has been idle, the less likely it is to return to production, although its observed state barely changes: capacity is frozen while a well is idle, and prices move little from one year to the next. The data cannot tell whether time idle itself makes a restart harder or whether the wells that remain idle are the ones that were always hard to restart. The structural model of Section 3 allows for the second: wells differ in a permanent reactivation cost, and a well can fail and lose its option to produce.

3 The structural model and its estimation

An operator decides each year whether to produce from a well (p), keep it idle (s, shut in) or plug it (e, exit). Plugging is permanent. The operator discounts the future at $\beta = 0.90$ per year. Wells differ in two ways we do not observe: a permanent type, which sets the cost of returning to production and scales the cost of plugging, and a failure state that the operator knows and we do not. Appendix D lists the notation.

3.1 State and transitions

The state of a well at the start of year t is the vector $x_t = (y_t, o_t, w_t, m_t)$. The status $y_t \in \{p, s\}$ is the action chosen the year before, a_{t-1}. Log production capacity o_t lies on a grid $\mathcal{O} = \{o_1, \ldots, o_{n_o}\}$ and the log real price w_t on a grid $\mathcal{W} = \{w_1, \ldots, w_{n_w}\}$, with $n_o = n_w = 31$. The failure state $m_t \in \{0, 1\}$ is 1 when the well has failed: it cannot produce without a repair. The operator knows m_t; we never observe it. Each well also carries a permanent type, introduced below.

Capacity follows a transition matrix Π_p^o for producing wells, estimated in a first stage from the panel and pooled over the wells of each fuel. We divide capacity into 31 bins whose limits are quantiles of the capacity of producing wells, with one grid node per bin, and assign each well-year to its bin. We then count, for every pair of bins, how often a well moves from one to the other between consecutive years; dividing each row by its total gives the empirical transition probabilities. To keep the matrix sparse, in each row we then set to zero the smallest entries whose probabilities add up to less than 5% and divide the remaining entries by their sum, so that the row again sums to one. Capacity does not depend on the price: once a well is drilled, its production responds little to prices (Anderson et al., 2018), so Π_p^o is the same at every price. An idle well's capacity does not change, $\Pi_s^o = I_{n_o}$. We do not observe recompletions, so a well that jumps to a higher capacity after a workover enters Π_p^o as an ordinary upward move, at no cost to the operator.

The price follows an AR(1) process common to all wells, $w_{t+1} = \mu_w(1 - \rho_w) + \rho_w w_t + b_w \eta_{t+1}^w$ with $\eta_{t+1}^w \sim N(0, 1)$, which we discretize on the grid into a transition matrix Π^w by the method of Tauchen

(1986). Since capacity and price move independently given the action, the exogenous state $\xi_t = (o_t, w_t)$ moves by the Kronecker product $\Pi_a^\xi = \Pi_a^o \otimes \Pi^w$ for $a \in \{p, s\}$: the probability of moving from (o, w) to (o', w') is $\Pi_a^o(o' \mid o)\,\Pi^w(w' \mid w)$. The model is solved on the grid. An observed well-year lies between nodes, and we evaluate its choice probabilities by bilinear interpolation between the four surrounding nodes (Section 3.4).

A producing well fails during the year with probability

$$\lambda(o) = \text{logit}^{-1}\big(\alpha_{\lambda 0} - \alpha_{\lambda 1}(o - \bar{o})\big), \tag{2}$$

evaluated at the current capacity o, where $\bar{o}$ is the mean of the capacity grid; $\alpha_{\lambda 1} > 0$ makes small wells fail more often. A failure happens while the well produces, and the operator learns of it at the start of the next year. To produce again, the operator must first pay the repair cost ψ_m; a repaired well that produces then faces the same failure hazard $\lambda(o)$ as any producing well. An idle well's failure state does not change. A well that enters the panel idle has failed with probability π_0^s, estimated; a well that enters producing has failed with probability $\lambda(o)$ at its first observed capacity.

3.2 Payoffs

Producing earns $u_p(o, w) = \theta_R \cdot W \cdot O - \theta_V \cdot O$. Here O is capacity in thousand barrels a year (thousand MMBtu for gas) and W the price in 2019 dollars per unit, so $W \cdot O$ is gross revenue in thousand 2019 dollars a year, and $\theta_V \cdot O$ is a cost per unit of capacity. Holding a well idle costs θ_0 a year, $u_s = -\theta_0$, which covers bonding, filings, surface maintenance and property tax. A plugged well earns nothing, $u_e = 0$. All costs are non-negative magnitudes that enter payoffs with a minus sign.

Changing status costs a one-time switching cost, with action first and origin second in the subscripts: plugging costs are denoted ψ_{ep} if plugged from a producing well and ψ_{es} from an idle one. Returning an idle well to production costs ψ_{ps}, and shutting in a producing well costs nothing. With rows for the current status (p, s) and columns for the action (p, s, e),

$$\Psi = \begin{pmatrix} 0 & 0 & \psi_{ep} \\ \psi_{ps} & 0 & \psi_{es} \end{pmatrix}, \tag{3}$$

plus ψ_m whenever a failed well produces. We allow the two plugging costs to differ because plugging an idle well can require more work: equipment and the wellbore deteriorate while the well sits idle, and the operator may need to restore the well's status and records before plugging, while a producing well has crews

and equipment on site.

The reactivation cost is the well's permanent type. Well i draws a score $z_i \sim N(0,1)$ once, when it enters the panel, and its reactivation cost is $\psi^i_{ps} = \exp(\mu_\psi + b_\psi z_i)$. The same score scales both plugging costs, $\psi^i_{ey} = \exp(\ln \psi_{ey} + \rho_e z_i) = \psi_{ey}\, e^{\rho_e z_i}$ for $y \in \{p, s\}$, so ψ_{ey} is the plugging cost of the median type and ρ_e plays for the plugging costs the role b_ψ plays for the reactivation cost. A positive ρ_e makes wells that are expensive to restart also expensive to plug, and $\rho_e = 0$ gives one plugging cost per status for all wells. The operator knows its type.

Each year the operator also draws a taste shock ε_a for each action, independent across actions, wells and years, Type I extreme value with mean zero. Its scale is one util. The operator does not receive all of the gross revenue $W \times O$: royalty owners and the state take a share s_R of it, so the operator's revenue is $(1 - s_R)\, W \times O$, and θ_R measures utils per thousand dollars of gross revenue. One util is therefore worth $(1 - s_R)/\theta_R$ of the operator's own (net) dollars, and every dollar figure in the paper is a parameter divided by θ_R and multiplied by $1 - s_R$. The model does not identify s_R, since it only rescales θ_R, so the estimates in utils do not depend on it. We set $1 - s_R = 0.76$, which corresponds to the average private royalty in Texas, about one fifth (Brown et al., 2016), and the Texas severance tax of 4.6% on oil, $0.80 \times 0.954 \approx 0.76$ (Texas Comptroller of Public Accounts, 2026a). The severance tax on gas is 7.5% (Texas Comptroller of Public Accounts, 2026b), which would give 0.74; we use one share for both fuels.

The cost of production is linear in capacity, and we restrict $\theta_V \geq 0$. The estimate sits at this bound in both fuels (Section 4), and every dollar figure in the paper is conditional on the restriction.

3.3 Timing and the Bellman equation

At the start of the year the operator sees the state vector $x_t = (y_t, o_t, w_t, m_t)$, its type and the shocks, chooses a_t, receives the flow payoff and pays any switching cost. Capacity and price then move by $\Pi^\xi_{a_t}$, and the failure state moves as described above. Write $\bar{V}(x)$ for the expected value of a well at state x before the shocks are drawn, with $\bar{V} = 0$ once plugged, and suppress the type. The produce continuation mixes over the failure draw,

$$\widetilde{EV}_p(o, w) = \sum_{o', w'} \Pi^\xi_p\big((o', w') \mid (o, w)\big)\Big[\big(1 - \lambda(o)\big)\, \bar{V}(p, o', w', 0) + \lambda(o)\, \bar{V}(p, o', w', 1)\Big].$$

The conditional values of the three actions at $x = (y, o, w, m)$, $y \in \{p, s\}$, are

$$\begin{aligned} v_p(x) &= u_p(o, w) - \psi_{ps}\,\mathbf{1}\{y = s\} - \psi_m\,\mathbf{1}\{m = 1\} + \beta\,\widetilde{EV}_p(o, w), \\ v_s(x) &= -\theta_0 + \beta\,EV_s(o, w, m), \qquad EV_s(o, w, m) = \sum_{w'} \Pi^w(w' \mid w)\,\bar{V}(s, o, w', m), \\ v_e(x) &= -\psi_{ey}, \end{aligned} \tag{4}$$

with ψ_{ps} and ψ_{ey} evaluated at the well's type. Because the shocks are Type I extreme value with mean zero,

$$\bar{V}(x) = \ln \sum_a e^{v_a(x)}, \qquad \sigma_a(x) = \exp\big(v_a(x) - \bar{V}(x)\big), \tag{5}$$

where $\sigma_a(x)$ is the conditional choice probability of action a; no Euler constant appears because the shocks have mean zero. We solve the fixed point by value function iteration, once for each type.

The expected shock of a well that chooses a is $\mathbb{E}[\varepsilon_a \mid a, x] = -\ln \sigma_a(x)$. A well of type score z observed to plug therefore bears $\psi_{ey} e^{\rho_e z} + \ln \sigma_e(x)$ in expectation, with σ_e evaluated at its type, which is less than its structural cost $\psi_{ey} e^{\rho_e z}$; Section 4 uses this to compute the realized plugging cost.

To keep the model to the decision of one well, we abstract from several features of the market. The model treats each well as its own decision maker, although on a multi-well lease the operator may decide for the lease as a whole. It has no operator portfolios, finance or bankruptcy, so it cannot say which wells become orphans, and no drilling or entry. It has one regulatory regime over 2001–2018, although HB 2259 changed the rules in 2010, and one type distribution for every well, whatever its age at entry. Holding a well idle costs the same in every year of a spell; Section 4.4 tests a cost that steps up with time idle.

3.4 Likelihood

We estimate the model by nested fixed point maximum likelihood (Rust, 1987), separately for oil and gas, on the estimation sample of Section 2.3, with the first-stage transitions Π^o_p and Π^w held fixed. The unobserved type is integrated out of each well's likelihood, as in the mixture models of Heckman and Singer (1984) and the dynamic discrete choice models with unobserved types of Keane and Wolpin (1997) and Arcidiacono and Miller (2011). The data for well i are its statuses, actions, capacities and prices over years $t = 1, \ldots, T_i$.

Fix the type at quadrature node q. The failure state is hidden, so the likelihood of the well's record comes from a forward filter over m.[13] Let $\hat{\alpha}_t(m)$ be the probability that the well is in failure state m at year t given its record before t, starting from $\hat{\alpha}_1(1) = \pi^s_0$ for a well that enters idle and $\lambda(o_{i1})$ for one that enters

[13] Because the choice probabilities are interpolated in logs (below), the interpolated probabilities of the three actions need not sum to one, so (6) is a composite likelihood rather than an exact probability of the record. It coincides with the exact likelihood at the grid nodes and as the grid is refined.

producing. The well's log-likelihood at node q is

$$\ell_{iq} = \sum_{t=1}^{T_i} \ln \sum_{m \in \{0,1\}} \hat{\alpha}_t(m)\, \tilde{\sigma}_{a_{it}}(x_{it}, q, m), \tag{6}$$

where the filter updates $\hat{\alpha}_t$ by Bayes' rule on the observed action and then moves it forward: unchanged after an idle year, and to $\big(1-\lambda(o_{it}), \lambda(o_{it})\big)$ after a producing year. The CCP $\tilde{\sigma}$ is evaluated at the observation's own price and capacity by bilinear interpolation of $\ln \sigma_a$ between the four surrounding grid nodes, in the spirit of the interpolation of Keane and Wolpin (1994), which avoids treating the grid as the data. The type is integrated out by quadrature: with Gauss–Legendre nodes u_q and weights ω_q, $q = 1, \ldots, 15$, the type score at node q is $z_q = \Phi^{-1}(u_q)$, and

$$\mathcal{L}(\theta) = \sum_{i=1}^{N} \ln \sum_{q=1}^{15} \omega_q\, e^{\ell_{iq}(\theta)}. \tag{7}$$

The model is solved once per node at each trial θ. We maximize $\mathcal{L}$ with a box-constrained quasi-Newton method, differentiating through the value function iteration and the filter by forward-mode automatic differentiation. Standard errors come from the inverse Hessian of $\mathcal{L}$; θ_V, which sits at its bound, is reported without one. The standard errors of dollar figures, which are ratios to θ_R, come from the delta method.

Table 5: Structural parameters of the estimated model

Parameter	Description	Restriction
θ_R	revenue coefficient (utils per thousand gross dollars)	> 0
θ_V	cost per unit of capacity	≥ 0 (binds)
θ_0	yearly cost of an idle well	≥ 0
ψ_{ep}, ψ_{es}	plugging cost from producing, from idle (median type)	≥ 0
μ_ψ, $\log b_\psi$	location and log-scale of the reactivation cost ψ_{ps}	—
ρ_e	loading of the plugging costs on the type score	—
$\alpha_{\lambda 0}$, $\alpha_{\lambda 1}$	failure hazard: logit level at $\bar{o}$, minus its slope in o	—
ψ_m	repair cost of a failed well	≥ 0
π_0^s	share of wells entering idle that have failed	$[0, 1]$

3.5 Identification

Appendix C gives the argument in two steps. The first recovers the choice probabilities of each type in each failure state from the panel. It is a finite-mixture problem with a hidden Markov state, and we rely on the results of Kasahara and Shimotsu (2009) and Hu and Shum (2012). These results need each well to be observed for several consecutive years, rank conditions that link the observed choices to the unobserved type and state, and enough variation in the observed state. Our type has a parametric (LogNormal) distribution, which the likelihood integrates over fifteen quadrature nodes; Appendix C discusses how the results apply to it. The second step inverts those probabilities type by type (Hotz and Miller, 1993), using log-odds against

plugging adjusted for the expected plug probability next year: the slope of the adjusted log-odds of producing in the price identifies θ_R, its level identifies the plugging cost from production, the gap between idling against plugging from the two statuses separates ψ_{ep} from ψ_{es}, idling against plugging from idle then gives θ_0, and, given the plugging costs, the gap between producing from a producing well and from an idle one gives each type's reactivation cost. The repair cost is the gap in the producing log-odds between a failed and a healthy producing well, and the LogNormal form and the loading ρ_e are over-identified across the fifteen types.

In the data, each part of the first step is tied to a different feature of the panel. The type distribution, μ_ψ and b_ψ, is identified by wells observed in more than one idle spell: a well that restarts quickly once tends to do so again, which a common reactivation cost cannot produce (Honoré, 1993). The loading ρ_e is identified by whether the wells that restart slowly also plug slowly. The failure state is identified by idle wells that respond less to prices: restarting a failed well also costs ψ_m, so a price rise moves it little. The estimate of ψ_m is large enough that a failed well almost never restarts: the data say that ψ_m is large, not how large. How the incidence of these wells varies with capacity identifies the slope $\alpha_{\lambda 1}$ of the hazard, and the wells that enter the panel idle pin down π_0^s. Therefore, in this model the falling hazard out of idleness comes from the type and the failure state, through dynamic selection (Heckman and Singer, 1984). A cost of idling that rises with time idle is excluded by assumption, and Section 4.4 checks it.

4 Estimates

4.1 Parameter estimates

Table 6 reports the twelve parameters of the reported specification in utility units, and Table 7 expresses them in thousands of net 2019 dollars, that is, divided by $\hat{\theta}_R$ and multiplied by $1 - s_R = 0.76$. One util is worth \$24.0k in oil and \$38.5k in gas. The variable-cost parameter θ_V sits at its lower bound of zero in both fuels. We keep this restriction, and every dollar figure is conditional on it.

The structural plugging cost of the median type is \$126.9k from a producing well and \$162.3k from an idle well in oil (\$144.0k and \$177.7k in gas). Plugging an idle well costs 28% more than plugging a producing one in oil and 23% more in gas, and the standard error of the difference is below 0.1 utils. These costs govern behavior in the counterfactuals; the costs actually paid by wells observed to plug are smaller, as Section 4.2 explains. Holding a well idle costs \$11.4k per year in oil and \$11.3k in gas, which covers bonding, the annual W-3X filing and extension conditions, surface maintenance and property tax. The two fuels agree to within the standard error.

A producing well fails with probability $\lambda(o)$ per year, 2.0% in oil and 1.0% in gas at the grid mean, and

Table 6: Structural parameter estimates, utility units

Parameter	Description	Oil	Gas
θ_R	revenue coefficient	0.0317***	0.0197***
		(0.0010)	(0.0005)
θ_V	variable/lifting cost	6.8e-14†	1.0e-17†
θ_0	idle flow cost	0.4732***	0.2928***
		(0.0063)	(0.0079)
ψ_{ep}	plugging cost, from producing	5.2884***	3.7368***
		(0.0649)	(0.0817)
ψ_{es}	plugging cost, from idle	6.7636***	4.6121***
		(0.0678)	(0.0828)
$\alpha_{\lambda 0}$	failure hazard, logit level at $\bar{o}$	-3.9141***	-4.5909***
		(0.0151)	(0.0285)
$\alpha_{\lambda 1}$	failure hazard, $-$slope in o	0.2564***	0.5927***
		(0.0040)	(0.0118)
ψ_m	repair cost	357.5979‡	278.2507‡
π_0^s	share of wells entering idle that have failed	0.3025***	0.1043***
		(0.0113)	(0.0193)
μ_ψ	ψ_{ps} mixing location	1.8795***	2.0003***
		(0.0074)	(0.0123)
$\log b_\psi$	ψ_{ps} mixing log-scale	-0.7809***	-0.6017***
		(0.0156)	(0.0212)
ρ_e	plugging-cost loading on the type score	0.0620***	0.0882***
		(0.0059)	(0.0107)
Log-likelihood		-343,226.2	-182,507.3
Free parameters		12	12
Well-years		1,071,027	598,760

Standard errors in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. † At its lower bound of zero; no standard error. All other parameters are interior. ‡ Large enough that a failed well does not produce; the data bound it from below only, since any larger value implies the same behavior, so no standard error is reported.

Table 7: Structural estimates in thousands of net 2019 dollars

Parameter	Description	Oil	Gas
Scale			
$(1-s_R)/\theta_R$	dollar value of one util	24.0***	38.5***
		(0.8)	(0.9)
Structural costs (median type)			
θ_V/θ_R	variable/lifting cost	1.6e-12†	3.9e-16†
θ_0/θ_R	idle flow cost	11.4***	11.3***
		(0.4)	(0.5)
ψ_{ep}/θ_R	plugging cost, from producing	126.9***	144.0***
		(4.4)	(5.8)
ψ_{es}/θ_R	plugging cost, from idle	162.3***	177.7***
		(5.4)	(6.5)
ψ_m/θ_R	repair cost	8,578.5‡	10,723.0‡
$e^{\mu_\psi}(1-s_R)/\theta_R$	median reactivation cost	157.1	284.8
Realized cost, conditional on switching			
$\check{\psi}_e$	plugging, all plug events	83.8***	59.6***
		(3.2)	(4.1)
$\check{\psi}_{ep}$	plugging from producing	62.8***	19.4***
		(2.8)	(3.7)
$\check{\psi}_{es}$	plugging from idle	93.7***	75.7***
		(3.4)	(4.4)
$\check{\psi}_{ps}$	reactivation	75.2***	109.7***
		(2.4)	(2.6)

Net of royalties and severance taxes, $1-s_R = 0.76$. Delta-method standard errors in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. † θ_V at its lower bound of zero. ‡ Large enough that a failed well does not produce; the data bound it from below only, since any larger value implies the same behavior, so no standard error is reported.

the hazard rises as capacity falls. A failed well must pay ψ_m on top of the reactivation cost to come back. The estimate is very large, $9 million in oil and $11 million in gas, so a failed well is effectively unrepairable. The failure state absorbs a recompletion decision the model does not contain: within the model, a well that would need a workover to produce again is permanently dead unless plugged. Of the wells that enter the panel already idle, 30% in oil and 10% in gas are estimated to be in the failed state. The plugging cost does not depend on the failure state; only the plug probability does.

The median reactivation cost is $157k in oil and $285k in gas, with log-scale $b_\psi = 0.46$ and 0.55, so the well at the 90th percentile of the type distribution faces a reactivation cost 1.8 times the median in oil and 2.0 times in gas. The same draw loads on the plugging costs through ρ_e, estimated at 0.062 (standard error 0.006) in oil and 0.088 (0.011) in gas. Wells that are expensive to restart are also somewhat more expensive to plug: a one-standard-deviation type moves the plugging cost by 6% in oil and 9% in gas.

4.2 The realized plugging cost

The structural plugging cost is the cost a well faces before it draws its year's shock, and the wells observed to plug drew a favorable one. For the mean-zero Type I extreme value shocks of the model, $\mathbb{E}[\varepsilon_a \mid a, q, m, x] = -\ln \sigma_a$, so the expected cost of plugging for a well of type q from status y that is observed to plug is $\psi_{ey} e^{\rho_e z_q} + \ln \sigma_e(x, q, m)$. The realized cost $\check{\psi}_e$ is its mean over the observed plug events, weighted by each well's smoothed posterior over its type and failure state. It is $84k in oil and $60k in gas. By status, plugs from producing wells cost $63k in oil and $19k in gas, and plugs from idle wells $94k and $76k. The same wells face a structural plugging cost, before their shock, of $151k in oil and $169k in gas on average, so the realized cost is 55% of the structural one in oil and 35% in gas; the rest is the favorable shock. Realized reactivation costs are $75k in oil and $110k in gas, below the median types' $157k and $285k, because the wells that reactivate are the cheaper types and seldom failed. The realized cost $\check{\psi}_e$ is the number to compare with engineering estimates, and the structural costs govern behavior in the counterfactuals.

Raimi et al. (2021) report a median decommissioning cost of $20k, rising to $76k when land reclamation is included, for orphan wells plugged by states; their figures are medians in their own dollar year; our realized cost is a mean in 2019 dollars. The model has no separate reclamation margin: the plugging cost is whatever the operator bears when the well leaves the data, so the relevant benchmark is the $76k figure that includes reclamation. Our oil estimate sits slightly above it and our gas estimate between the two figures.

The realized cost is a dollar amount: what the average well that plugs pays to plug, in net 2019 dollars, its shock included. The model also bounds it from above. Instead of plugging, an operator can always keep the well idle forever, at a present cost of $\theta_0/(1-\beta)$, which is $114k in oil and $113k in gas. An operator

who pays more than that to plug would have done better by idling, so the realized cost of the wells that plug lies below this amount, which the idle cost and the discount factor set (Appendix C, equation (11)). A lower benchmark is the state's own spending: the Railroad Commission spent $17,012 per well it plugged in fiscal 2015 (SAC, 2016), which covers plugging only and excludes site remediation. Contractor costs for state-selected orphans exclude the operator's own reclamation, filings and administration, all of which our estimate can contain, so this comparison bounds the estimate from below rather than testing it.

The homogeneous model, which has neither the type nor the failure state (column (1) of Table 8), helps explain why the heterogeneity matters for this number. Without types and the failure state, one plugging cost must rationalize both the wells that plug and the long-idle wells that never do. The homogeneous model's structural plugging cost is therefore large, $378k in oil and $318k in gas, and its plug probabilities so low that the favorable shock of a well observed to plug almost offsets the cost: its realized cost is $40k in oil and −$11k in gas (Table 8). The negative figure diagnoses the homogeneous model and is not a cost estimate.

4.3 Model fit

Figures 5–8 compare the model's one-step-ahead action probabilities with the observed action shares in cells of price quintile, capacity quintile, status and years idle, resetting each well to its observed state each year. The model captures the broad patterns: continued production rises with capacity and plugging is more likely at low prices. For oil wells that enter the year idle, the model under-predicts the probability of remaining idle and over-predicts the return to production. The oil plug-from-idle share is humped in capacity, 4.8, 6.5, 5.6 and 4.5% from the second capacity quintile up, while the model predicts 5.4, 5.0, 4.8 and 4.9%. The hump is in the observed shares, and it is absent on single-well leases (Section 4.4). For gas wells that enter the year producing, the model over-predicts continued production at the lowest capacity quintile and under-predicts shutting in.

Figure 9 is the stronger test. From each well's observed entry state and under the realized price path, the fitted model is propagated forward and compared with the data year by year, each year carrying its own simulated survivors. The model reproduces the level of plugs and of the idle share through about 2010 and drifts above both afterwards. It over-predicts total plugs by 8% in oil and 10% in gas, and the idle share of the 2018 stock by 4.5 and 5.8 percentage points. Two misses are visible. The first follows the RRC's implementation of HB 2259 in September 2010, which attached a financial-assurance requirement to every plugging extension (Section 2). In oil, plugs jump in 2011–12 and the model does not produce the jump. In gas, plugs rise to a peak in 2013 and fall sharply after 2014; the model follows the rise, through its drift, but not the fall. The regime change is not in the model. The second miss is the gradual over-prediction of plugs,

Figure 5: Empirical (black) and model (blue) choice probabilities for oil wells entering the year producing, by price and capacity quintile

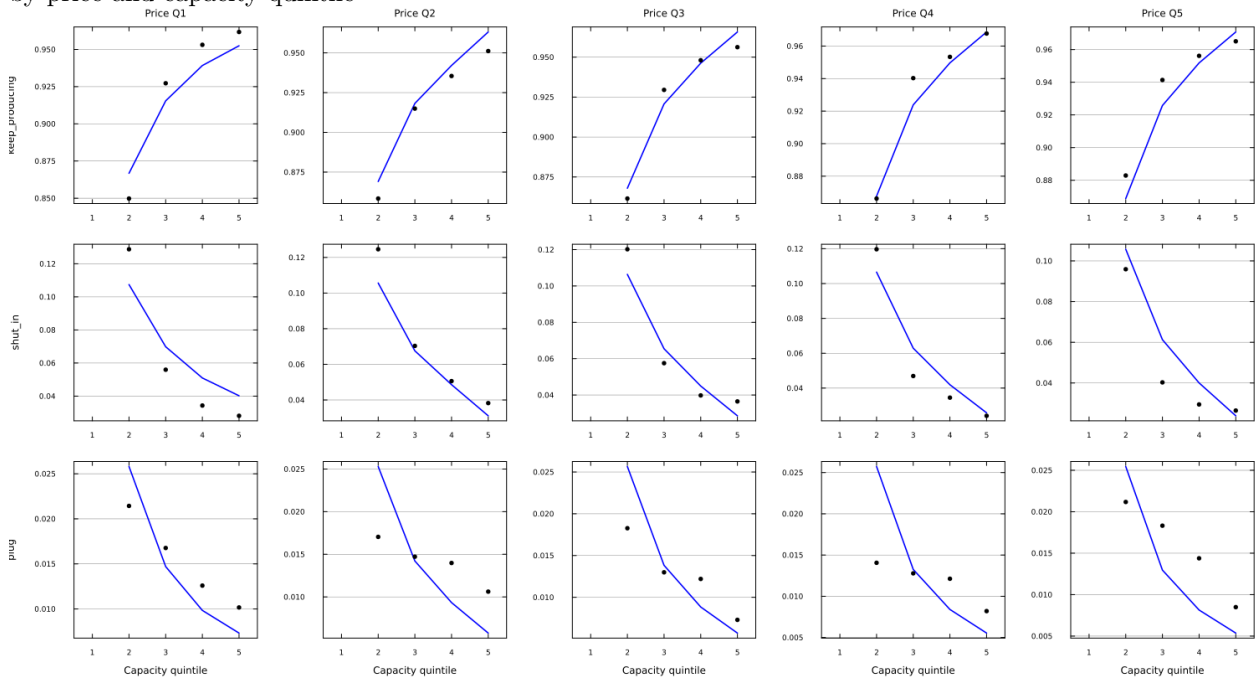

Figure 6: Empirical (black) and model (blue) choice probabilities for oil wells entering the year idle, by price and capacity quintile

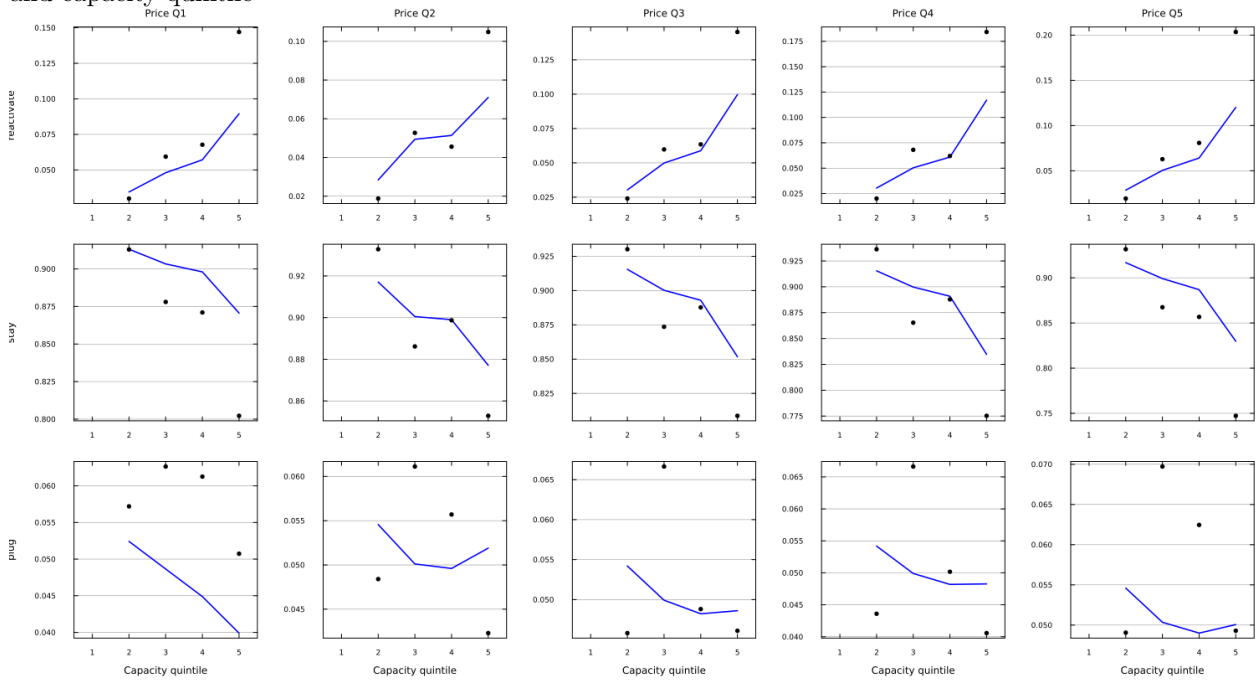

Figure 7: Empirical (black) and model (blue) choice probabilities for gas wells entering the year producing, by price and capacity quintile

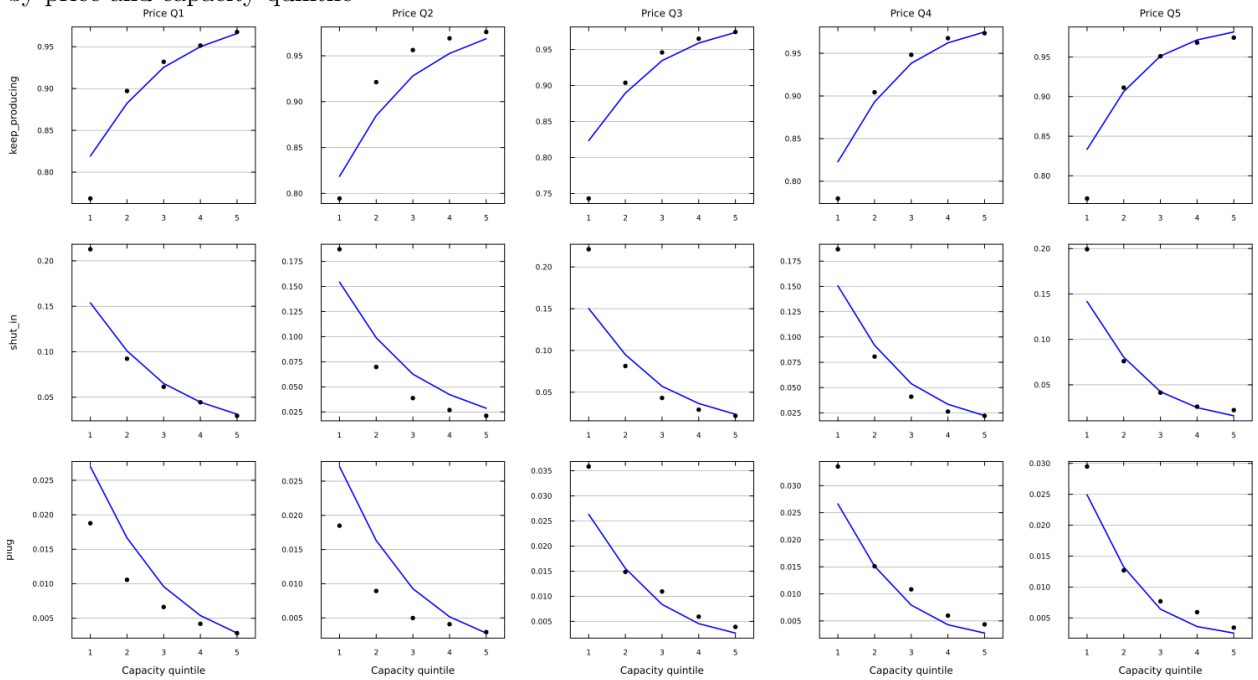

Figure 8: Empirical (black) and model (blue) choice probabilities for gas wells entering the year idle, by price and capacity quintile

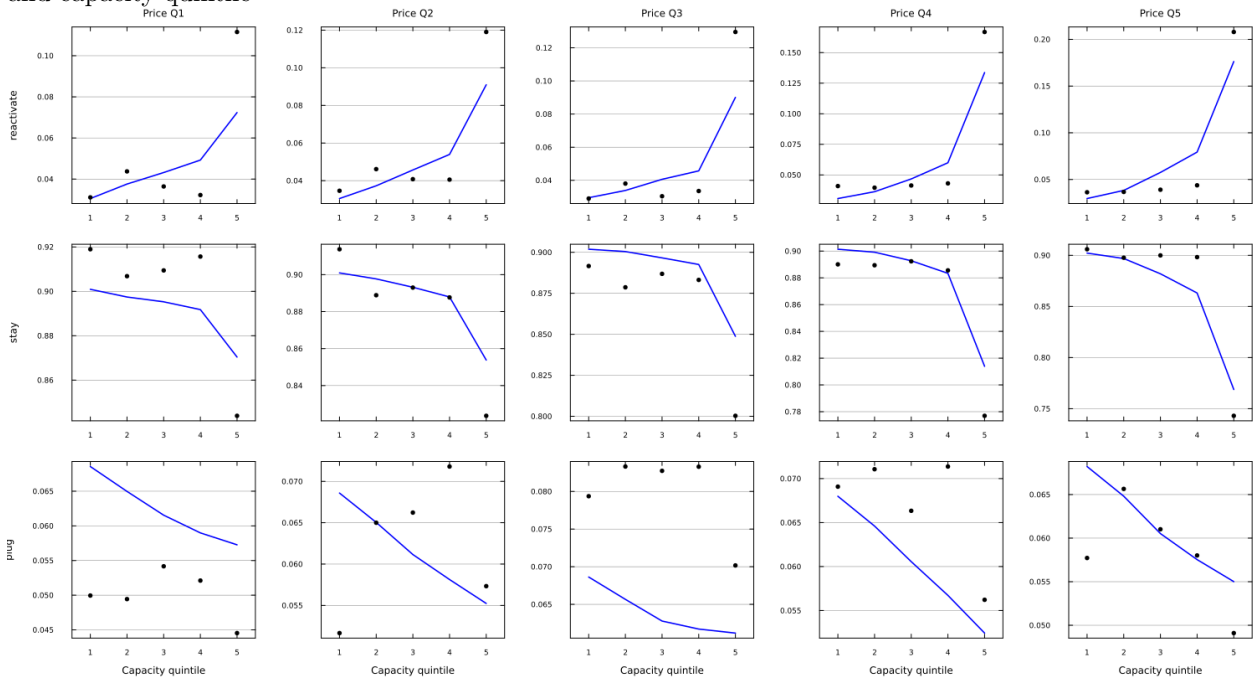

which comes from the first-stage capacity transition. The trim of Section 3.1 drops, at low capacity, mostly the upward moves, so the simulated stock sinks toward the bottom of the grid, where plugging hazards are highest, faster than the data do. A matrix without the trim, with the reported model re-estimated on it, brings the plug count to within 3.5% of the data in oil and 1.3% in gas. We keep the stored matrix and report the drift.

Figure 9: Stage 1 backcast. From each well's entry state and under the realized prices, the fitted model's plugs and idle share against the observed panel, 2001–2018

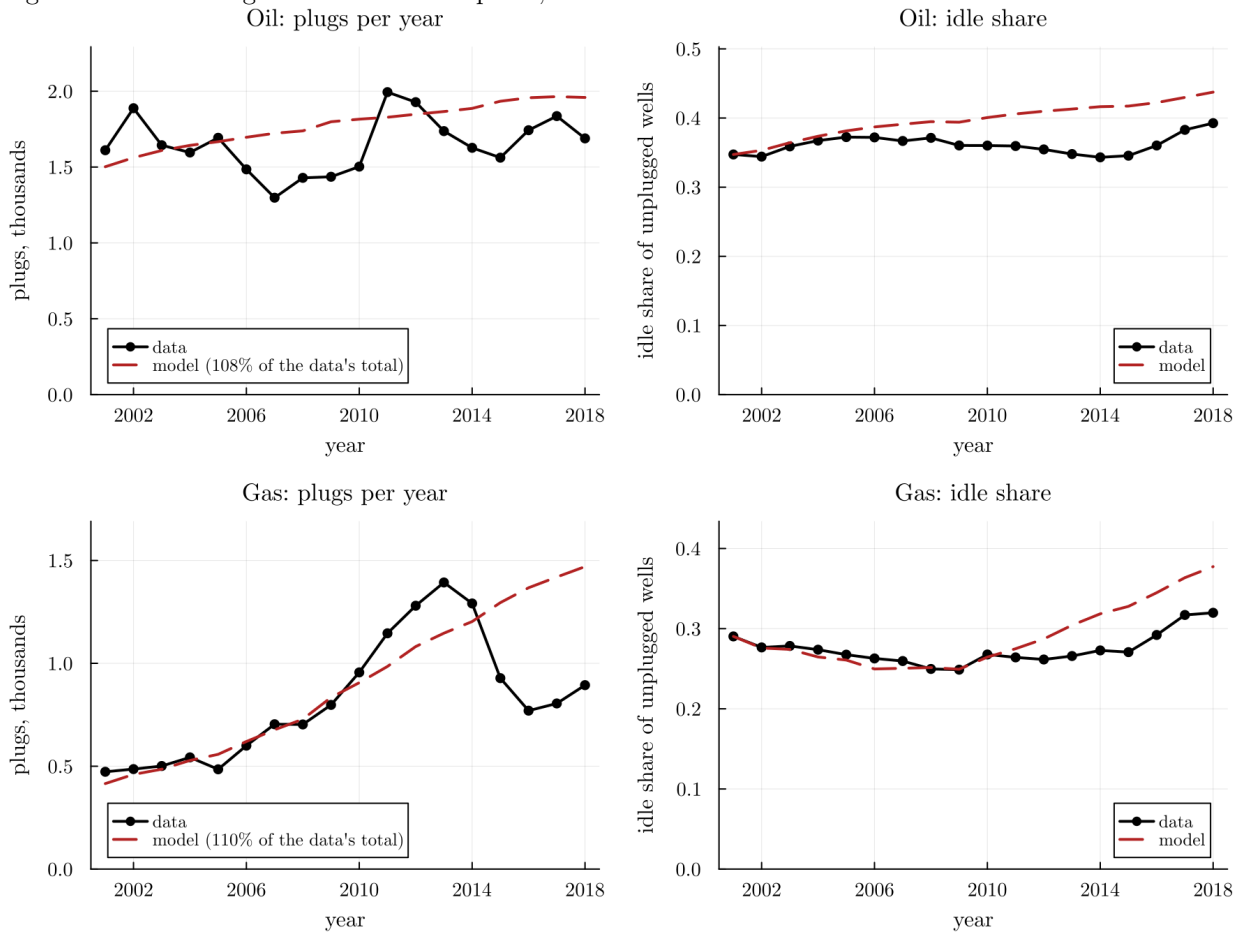

4.4 Specification comparison and robustness

Table 8 reports the ladder from the homogeneous model, in column (1), to the reported one, in column (4), on the estimation sample; the bottom of the table ticks the features of each column. Adding the hidden failure state with status-dependent plugging costs (column 2) gains 13,739 (oil) and 5,996 (gas) log-likelihood points for five parameters. Adding the random reactivation cost (column 3) gains a further 2,044 and 1,077 for one, and the type loading on the plugging cost, ρ_e (column 4), a further 62 and 36 for one. The Bayesian

information criterion selects the reported specification over everything below it on the ladder. Along the ladder the plugging cost falls from \$378k in the homogeneous model to \$126.9k in the reported one (oil; gas \$318k to \$144.0k), and the realized cost rises from \$40k to \$84k. The homogeneous model needs a large structural cost to rationalize the low plug rate of wells that, in the richer model, are cheap-to-restart types or failed wells waiting for a favorable shock.

Table 8: Specification comparison on the estimation sample

	(1)	(2)	(3)	(4)	(5)	(6)
Oil						
Log-likelihood	-359,072	-345,333	-343,288	-343,226	-343,195	-343,267
Free parameters	5	10	11	12	13	12
BIC	718,213	690,804	686,730	686,619	686,571	686,701
ψ_{ep} (k\$)	377.5	147.1	135.0	126.9	135.2	136.2
ψ_{es} (k\$)	—	185.5	172.5	162.3	173.9	178.6
θ_0 (k\$/yr)	23.2	13.2	12.1	11.4	12.9	9.9
realized $\tilde{\psi}_e$ (k\$)	40.0	100.5	92.5	83.8	91.8	97.4
Gas						
Log-likelihood	-189,617	-183,621	-182,544	-182,507	-182,489	-182,542
Free parameters	5	10	11	12	13	12
BIC	379,300	367,376	365,234	365,174	365,151	365,244
ψ_{ep} (k\$)	318.0	163.8	149.2	144.0	159.2	149.0
ψ_{es} (k\$)	—	199.3	184.5	177.7	195.3	186.6
θ_0 (k\$/yr)	18.7	13.1	11.7	11.3	13.5	10.6
realized $\tilde{\psi}_e$ (k\$)	-10.9	75.8	66.8	59.6	73.1	68.6
Features						
hidden failure state		✓	✓	✓	✓	✓
$\psi_{ep} \neq \psi_{es}$		✓	✓	✓	✓	✓
random reactivation cost ψ_{ps}			✓	✓	✓	✓
type loading on plugging, ρ_e				✓	✓	
plugging premium for failed wells					✓	
step in the idle cost after year 3						✓

Column (4) is the reported specification. Dollar rows in thousands of net 2019 dollars. In every column θ_V is at its lower bound of zero; all other parameters are interior.

Two variants extend the reported model. The first adds a plugging premium for failed wells (column 5). It raises the log-likelihood by 31 points in oil and 18 in gas for one parameter, and the Bayesian information criterion prefers it in both fuels.

The second adds a step in the idle flow cost: from the third consecutive idle year on, holding the well idle costs $\theta_0 + \theta_I$ a year instead of θ_0 (column 6, added to the model without ρ_e). The step raises the yearly holding cost by 28% in oil, which is significant ($z = 6.5$, 21.5 log-likelihood points), and by 13% in gas, which is not ($z = 1.7$). The type and the failure state remain in the model with the step, so the falling hazard out of idleness is mostly dynamic selection, with a modest rise in the holding cost in oil. We keep the specification without the step: the step is small in oil and not significant in gas. Leaving out a step that is truly there would bias the remaining estimates, mostly the idle cost and the type distribution in oil, while adding one

that is not there would only cost precision.

The last check re-estimates oil on the single-well leases of the estimation sample, where the lease volume is the well's own and the Enverus allocation cannot enter. There are 13,501 such oil wells, 14.8% of the oil sample, with 134,413 well-years. On this subsample the plugging cost is \$50k (standard error 4.5) from a producing well and \$71k (5.2) from an idle one, the idle flow cost \$4.3k per year (0.4), the type loading on the plugging cost 0.36 (0.03), and one util is worth \$25.3k, against \$24.0k in the pooled oil fit. The realized plugging cost is \$-0.8k, about zero. The plug-from-idle share is flat in capacity here, so the hump of the model fit is absent. Two readings are possible. Either, on multi-well leases, the allocation keeps dead wells "producing" and inflates the pooled plugging cost, or single-well leases are a selected population, with smaller operators and older fields, whose wells are cheaper to plug. The data we have cannot separate the two. We keep the pooled estimates as the main set because they cover the whole population the policies apply to, and Section 5 shows that the instrument ranking is the same on single-well leases.

4.5 How much the heterogeneity matters

Three facts measure how much the heterogeneity matters. The first is the likelihood: the reported model gains 15,845 log-likelihood points in oil and 7,109 in gas over the homogeneous model for 7 additional parameters, and the Bayesian information criterion falls from 718,213 to 686,619 in oil and from 379,300 to 365,174 in gas.

The second is the fit shown in Figures 5–8. For each cell of price quintile, capacity quintile, status and years idle, we take the absolute difference between the predicted and the observed share of an action, and we average these differences across cells, weighting each cell by its number of well-years. For plugging from idle in oil, this average error falls from 0.020 in the homogeneous model to 0.013 in the reported one, and for shutting in from 0.026 to 0.011, where 0.010 is one percentage point (gas: 0.020 to 0.016 and 0.030 to 0.014).

The third is the one the counterfactuals depend on. Take observably identical idle wells, at the median price and the median capacity of the oil sample. The homogeneous model gives each of them the same plug probability, 6.8% a year. The reported model gives each type its own probability: across the fifteen types it ranges from 0.31% to 5.4% for a healthy well, with a weighted mean of 3.7%, and it is 6.7% for a failed well. The homogeneous model's probability is above every healthy type's at the median, so it treats the average idle well as readier to plug than any healthy type is. At every capacity it lies above the type-weighted mean: at the lowest capacity the healthy types range from 0.63% to 5.7% against the homogeneous model's 7.0%, and at the 90th percentile of capacity from 0.03% to 4.3% against 5.9% (gas: 0.89 to 7.7 against 8.1, and 0.02 to 6.0 against 5.5). Therefore, a regulator who sets an instrument from the homogeneous model treats

every idle well as the average well, and Section 5.6 shows what that costs.

5 Counterfactuals

We compare four instruments that a regulator could use to bring plugging forward, and add a fifth, a stylized version of the bond Texas requires. For each we re-solve the estimated model under the modified payoffs or action set and follow the 2018 stock of unplugged wells forward. All results are for the low-producing wells of Section 2.3.

5.1 Instruments and population

A plugging subsidy ς is paid when a well is plugged from either status. It lowers every type's plugging cost by the same ς dollars, that is, by $\theta_R\varsigma/(1-s_R)$ utils, whatever the type's loading $e^{\rho_e z}$; we solve it on a grid from \$2.5k to \$500k per plug. An idling tax τ_s is charged for every year a well is idle and raises the yearly idle cost by τ_s dollars, from θ_0 to $\theta_0 + \theta_R\tau_s/(1-s_R)$ utils, on a grid from \$0.25k to \$50k per idle year. A capacity mandate $\bar{O}$ forces the plugging of any well, producing or idle, whose production capacity falls to or below $\bar{O}$; a well above the threshold is free to produce, idle or plug. We solve it at each of the 31 nodes of the capacity grid. Our headline threshold is 3.3 barrels a month in oil and 196 MMBtu a month in gas, so wells at the four lowest capacity nodes of the grid must plug. Texas does not use such a rule; we include it as a standard that acts on how much a well can produce rather than on how long it has been idle.

An idling limit $\bar{\kappa}$ removes the idle action after $\bar{\kappa}$ consecutive idle years, so a well at the limit must reactivate or plug. It is equivalent to enforcing the inactive-well rule without the low-cost extensions that let operators idle indefinitely (Section 2). Our headline is $\bar{\kappa} = 2$, the yearly counterpart of the twelve-month rule checked at each annual renewal, and we report $\bar{\kappa} = 1$, 3 and 5 alongside. The estimated model has no count of consecutive idle years, so for the limit we add one to the state; it does not enter payoffs, so without the limit the model is unchanged.

The bond is a uniform yearly cost τ_b on every unplugged well, producing or idle. For a letter of credit, τ_b is the premium rate times the bond amount: at a premium of 5% a year, a bond of \$10k per well costs \$0.5k a year. A higher bond amount or a higher premium rate raises τ_b in proportion. In behavior the bond equals a plugging subsidy of $\tau_b/(1-\beta)$. A well that plugs at date T pays $\tau_b(1-\beta^T)/(1-\beta)$ in present value, which is the constant $\tau_b/(1-\beta)$ minus a subsidy of $\tau_b/(1-\beta)$ received at T, and the constant does not affect any decision. Unlike the subsidy, however, the bond's payment is a pure cost to the operator that no one in this market receives, so all of it counts as a loss. We use τ_b from \$0.25k to \$1k per well-year.

The population is every well of the estimation sample still unplugged after its 2018 decision: 61,281 oil

wells and 45,289 gas wells. Each starts from its posterior over the reactivation-cost type and the failure state given its 2001–2018 record, and prices follow the estimated AR(1) from the 2018 node. We abstract from drilling, from entry into the low-producing group from above, and from operators' portfolios. The regulator sees price, capacity, status and years idle, and knows the estimated distribution of types and failure states, but not a given well's type or whether it has failed.

Under each policy the model is re-solved at each of the 15 type nodes. We then compute, without simulating individual wells, how the share of wells in each state evolves: starting from the 2018 distribution, each year's distribution is the previous one multiplied by the policy's transition probabilities. Adding up the discounted distributions over all future years gives, for every state, the expected discounted number of well-years spent there, which we obtain at once by solving a linear system. From these we read off the discounted idle well-years $\mathcal{I}$, producing well-years $\mathcal{Y}$ and plugs, the operators' value $\mathcal{U}$ and the government's revenue G. Run from each well's entry state, this calculation reproduces the 2001–2018 data up to the drift shown in Figure 9, and it agrees with a Monte Carlo simulation of the same policies.

5.2 Measuring abatement and welfare

Every result compares a policy with a baseline: the same 2018 stock followed forward under the rules in force during our sample, extensions included, as the estimated costs embed them. "No policy" below means this baseline. In the baseline the oil stock spends 224 thousand discounted idle well-years, and the gas stock 154 thousand. We measure abatement in discounted well-years. An idle well-year avoided next year counts as $\beta = 0.9$, one avoided in two years as 0.81, and so on, so $\Delta\mathcal{I}$, the baseline value of $\mathcal{I}$ minus its value under the policy, is the discounted number of idle well-years the policy avoids, and $\Delta\mathcal{I}$ divided by the baseline $\mathcal{I}$ is the share of the baseline's idle well-years it avoids. The same holds for producing well-years, $\Delta\mathcal{Y}$. Damages are linear in well-years: d_s dollars per idle well-year and $d_p = \chi d_s$ per producing well-year. A policy's abatement in idle-year equivalents is then

$$\Delta\mathcal{B}_\chi = \Delta\mathcal{I} + \chi\,\Delta\mathcal{Y}, \tag{8}$$

and its benefit is $d_s\,\Delta\mathcal{B}_\chi$. When $\chi = 0$ only idle years count; when $\chi = 1$ a producing year and an idle year count the same, so $\Delta\mathcal{B}_1$ is the discounted number of unplugged well-years avoided. The deadweight loss of the policy is the fall in $\mathcal{U} + G$, the operators' value plus the government's revenue, relative to the baseline, in net dollars. It counts no damage: it is what the policy costs operators, net of the transfers between them and the government, and the damage avoided is counted separately, as $d_s\,\Delta\mathcal{B}_\chi$. The net benefit of the policy is $d_s\,\Delta\mathcal{B}_\chi - \mathrm{DWL}$.

A regulator chooses an instrument and its level, not an abatement. Each level of an instrument, a tax of

so many dollars per idle year or a limit of so many years, gives one abatement, and a higher level gives more. We can therefore line up instruments by the abatement they achieve: each limit and each mandate threshold gives one abatement, and for that abatement there is exactly one tax and one subsidy that achieve it. The figures below put the abatement on the horizontal axis for this reason. A vertical line through a figure picks out one abatement, and the points where it crosses the curves are the levels of each instrument that achieve it.

Section 5.3 uses this alignment to rank instruments by deadweight loss at equal abatement. Section 5.4 fixes the policy and asks how large the damage must be for it to raise welfare. The net benefit is positive if and only if

$$d_s > D^* = \frac{\text{DWL}}{\Delta \mathcal{B}_\chi}, \tag{9}$$

so D^*, the deadweight loss per idle-year equivalent avoided, is the break-even damage per idle well-year: the external cost an idle well-year must do for the policy to pay. It needs no damage estimate, so a reader with their own damage estimate can read off which policies pass. In Section 5.5 we instead fix the damage per idle well-year and ask what each instrument gains. The damage can come from any source, such as methane, groundwater contamination, lost property value or the fiscal cost of orphaned wells, but only methane has a value in the literature, so we use it as the benchmark. The methane damage per idle well-year is the leak rate times the social cost of a gram per hour for a year: at 15 g/h and \$16.2 per g/h-year (Table 16 in Appendix B lists the measured leak rates), \$243, which we round to $d_s = \$0.25$k.

The benchmark is the first-best fee: a fee equal to the damage on every damaging well-year, d_s per idle year at $\chi = 0$, and $(\tau_s, \tau_p) = (d_s, \chi d_s)$ on idle and producing years at $\chi > 0$. The fee prices the externality where it occurs and leaves every other margin to the operator, so no instrument abates as much at a lower deadweight loss. Throughout, "the fee" means a fee of this form at whatever level achieves a given abatement, the cheapest way to achieve it; at $\chi = 0$ it is the idling tax, which the tables label "idling tax (the fee)", and every deadweight-loss ratio below divides by its loss. Its level is first-best only when it equals the damage. Our headline is $\chi = 0$, with $\chi = 1$ alongside. The question of this paper is the idle well, and most of the methane of a producing year comes with the oil it produces, which in part replaces oil produced elsewhere with its own emissions, so the leak rates of producing wells in Table 16 overstate the damage a restarted well adds.

5.3 The same abatement at different costs

Consider a regulator who enforces an idling limit of 2 years on the oil wells. Under the limit, the 2018 stock spends 210 thousand fewer discounted idle well-years than in the baseline, a 94% decrease from the 224

thousand it would otherwise spend idle (in gas, 139 thousand, 91%). Because the limit removes almost all idle years, a price that matches it must be large: a regulator who wanted the same reduction from a price would set an idling tax of $30.76k per idle year or a plugging subsidy of $331k per plug. Table 9 compares the three, at $\chi = 0$, and does the same for the capacity mandate at its headline threshold. That tax has a deadweight loss of $1.50 billion; in gas the tax is $40.19k and its loss $1.48 billion. The limit's own loss is $1.94 billion in oil and $1.99 billion in gas, which is 1.29 and 1.35 times the tax's. A plugging subsidy that achieves the same reduction, $331k per plug in oil and $465k in gas, has a loss of $4.13 billion in oil and $4.39 billion in gas, 2.75 and 2.97 times the tax's.

The instruments reach the same reduction in different ways. The subsidy abates by inducing plugs, including of wells that would have returned to production, so it forgoes 179 thousand discounted producing well-years in oil. The tax and the limit instead make idling costly or impossible, so some idle wells plug and others return to production: they add 25 and 22 thousand producing well-years, the negative entries in the producing column of Table 9. The capacity mandate at its headline threshold abates less than the limit, so the second panel of Table 9 sets the tax and the subsidy to the mandate's own reduction, 129 thousand idle well-years in oil: each comparison in the table is between policies that achieve the same thing. The mandate costs 9.2 times the tax that achieves it (9.8 in gas), because it forces the plugging of producing wells with years of production ahead of them (136 thousand producing well-years forgone in oil). As the simple model of Figure 1 suggests, the limit and the mandate lose because they cannot tell wells apart: a limit forces every well idle for $\bar{\kappa}$ years to decide, including wells whose option to restart is valuable.

The tighter the limit, the more it abates and the closer it comes to the fee (Table 10). The one-year limit costs 1.13 times the fee at its abatement in oil, the three-year limit 1.45 times and the five-year limit 1.79 times. A loose limit acts on few wells, those idle for many years, while the fee still selects, among all idle wells, those with the least to give up.

Figure 10 draws the same comparison for every level of every instrument. Each point is one policy: its horizontal position is the abatement the policy achieves, in discounted idle well-years avoided, and its height is the policy's deadweight loss, on a log scale. The tax and the subsidy are curves, because their level can be set continuously. Moving right along the tax curve raises the tax, from $0.25k per idle year at its left end to $50k at its right end; moving along the subsidy curve raises the subsidy from $2.5k to $500k per plug. These are levels of the instruments, not damages; no damage enters the figure. Each idling limit and each mandate threshold is a single point. Read the figure along a vertical line. The dotted line in panel (a) is at the 210 thousand idle well-years the limit of 2 years avoids in oil. The limit is the point on it at a height of $1.94 billion (the axis is in millions). The line crosses the tax curve at the tax of $30.76k per idle year, at $1.50 billion, and the subsidy curve at the subsidy of $331k per plug, at $4.13 billion: these are the three rows of

Table 9: Instruments compared at equal abatement of idle well-years ($\chi = 0$)

	Level (k$)	Idle well-years avoided (k)	Producing well-years avoided (k)	DWL ($ million)	DWL relative to the fee
Oil					
Same abatement as the idling limit $\bar{\kappa} = 2$					
idling limit	2 years	210.1	-22.5	1,937.2	1.29
idling tax (the fee)	30.76 per idle year	210.1	-24.8	1,504.6	1.00
plugging subsidy	331 per plug	210.1	179.0	4,131.8	2.75
Same abatement as the capacity mandate					
capacity mandate		129.2	136.0	3,468.0	9.23
idling tax (the fee)	6.86 per idle year	129.2	-15.8	375.9	1.00
plugging subsidy	77 per plug	129.2	50.2	642.4	1.71
Gas					
Same abatement as the idling limit $\bar{\kappa} = 2$					
idling limit	2 years	139.4	-19.9	1,991.5	1.35
idling tax (the fee)	40.19 per idle year	139.4	-23.8	1,477.4	1.00
plugging subsidy	465 per plug	139.4	131.9	4,393.6	2.97
Same abatement as the capacity mandate					
capacity mandate		91.9	113.1	4,442.7	9.80
idling tax (the fee)	11.92 per idle year	91.9	-15.6	453.3	1.00
plugging subsidy	142 per plug	91.9	49.0	913.1	2.01

Net 2019 dollars; discounted well-years ($\beta = 0.90$), thousands. A negative number of producing well-years avoided means the policy returns idle wells to production.

Table 10: Idling limits of one to five years ($\chi = 0$)

Idling limit	Idle well-years avoided (k)	Producing well-years avoided (k)	DWL relative to the fee		Break-even damage D^* (k$/idle yr)
			limit	subsidy	
Oil					
$\bar{\kappa} = 1$	219.2	-24.9	1.13	2.97	9.57
$\bar{\kappa} = 2$ (headline)	210.1	-22.5	1.29	2.75	9.22
$\bar{\kappa} = 3$	199.4	-19.9	1.45	2.53	8.94
$\bar{\kappa} = 5$	176.4	-15.1	1.79	2.16	8.45
Gas					
$\bar{\kappa} = 1$	148.0	-23.2	1.17	3.24	15.09
$\bar{\kappa} = 2$ (headline)	139.4	-19.9	1.35	2.97	14.28
$\bar{\kappa} = 3$	130.0	-16.9	1.53	2.71	13.66
$\bar{\kappa} = 5$	110.7	-12.1	1.91	2.29	12.57

The subsidy and the fee are set to the idle well-years the limit avoids. Net 2019 dollars.

Table 9. At any abatement, the lower an instrument lies, the cheaper it is, and every idling limit and every mandate threshold lies above the tax curve. Even the fee has a deadweight loss, because forcing operators to plug or restart earlier than they would have costs them; the benefit, the damage avoided, is counted in the next subsection.

Figure 10: Deadweight loss against idle well-years avoided, by instrument, $\chi = 0$: (a) oil, (b) gas. The curves trace the idling tax ($0.25k to $50k per idle year) and the plugging subsidy ($2.5k to $500k per plug) as their level rises from left to right; the points are the idling limits of 1 to 12 years and the 31 mandate thresholds. The dotted vertical line is at the abatement of the limit of 2 years, and the labelled markers are the tax and the subsidy that achieve the same abatement

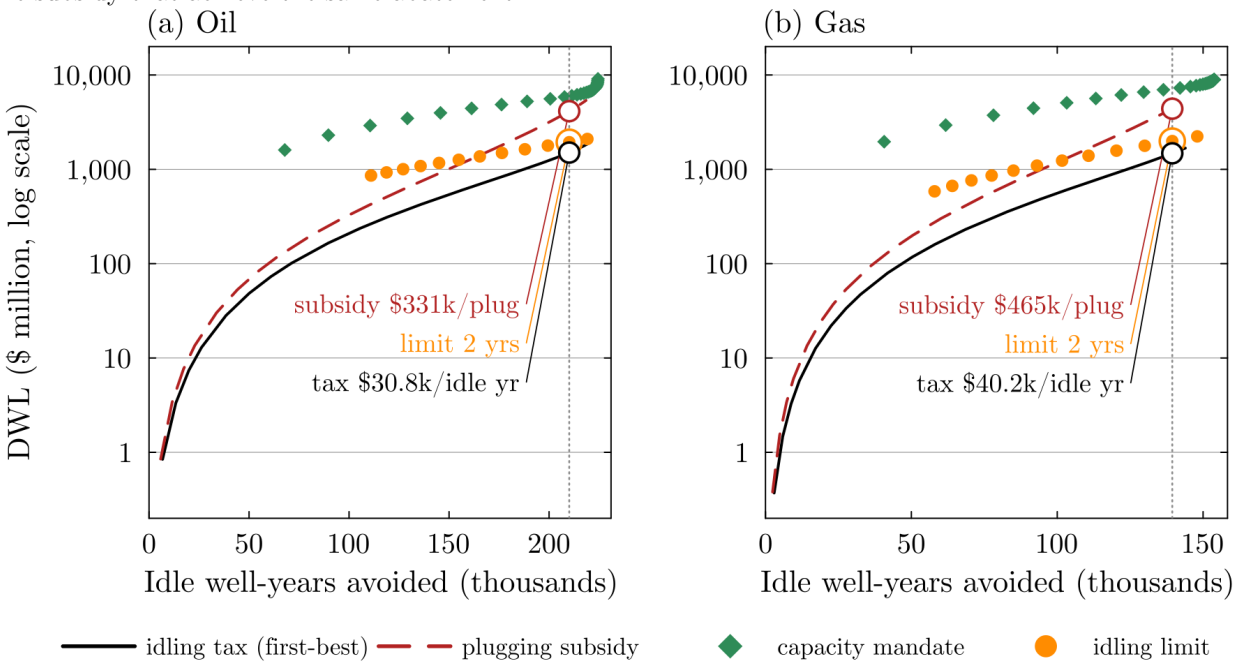

When producing wells also do damage, the ranking of the price instruments changes (Table 11). Take the limit of 2 years and suppose a producing year does damage χ times that of an idle year. For each χ we find the two-part fee that delivers the same abatement $\Delta\mathcal{B}_\chi$ as the limit; it is the cheapest way to do so, and we divide each instrument's deadweight loss by the fee's. At $\chi = 0$ the fee is the idling tax alone, so the tax's entry is 1. At $\chi = 1$, when idle and producing years do equal damage, the subsidy matches the fee (ratio 1.00 in oil and 1.00 in gas): a fee charged equally on every unplugged year behaves as a subsidy paid at plugging. An idling tax alone, set to the same $\Delta\mathcal{B}_1$, then costs 2.2 times the fee in oil and 2.6 times in gas, and the limit 2.7 and 3.1 times, because both return idle wells to production, which now counts as damage. The more damage producing years do, the more the limit and the tax fall behind. The limit stops abating altogether once χ exceeds $\Delta\mathcal{I}/|\Delta\mathcal{Y}|$, 9.3 in oil and 7.0 in gas, because the producing years it adds outweigh the idle years it avoids; the tax that matches the limit's reduction in idle well-years stops abating above 8.5

and 5.9.

Table 11: Instruments at the abatement of the idling limit, by the damage ratio χ

$\chi = d_p/d_s$	Abatement $\Delta\mathcal{B}_\chi$ (k well-years)	DWL relative to the fee		
		limit $\bar{\kappa} = 2$	idling tax	subsidy
Oil				
0	210.1	1.29	1.00	2.75
0.3	203.3	1.49	1.17	1.28
1	187.6	2.74	2.25	1.00
3	142.7	15.29	15.40	1.51
5	97.7	77.81	–†	2.45
Gas				
0	139.4	1.35	1.00	2.97
0.3	133.5	1.55	1.19	1.32
1	119.6	3.11	2.65	1.00
3	79.8	25.91	–†	1.56
5	40.0	263.12	–†	2.53

† The tax cannot reach the target within its grid (up to \$50k per idle year).

5.4 How large must the damage be?

Equation (9) turns each policy into one number, the damage per idle well-year it needs to pay for itself. No damage value enters it: D^* is the policy's deadweight loss divided by its abatement $\Delta\mathcal{B}_\chi$, which is the idle well-years it avoids when $\chi = 0$, so it is in dollars per idle-year equivalent avoided. A policy raises welfare if an idle well-year does more damage than D^*, from whatever source. We compare D^* with the methane damage of \$0.25k only because methane is the one damage the literature values. For the limit of 2 years, D^* is \$9.2k in oil (\$14.3k in gas). That is, the limit raises welfare only if an idle well-year does at least \$9.2k of external damage in oil, through every channel together. At the same abatement, D^* is \$7.2k in oil (\$10.6k in gas) for the idling tax, \$19.7k (\$31.5k) for the subsidy, and \$27k (\$48k) for the mandate at its own abatement. Methane is about \$0.25k per idle well-year, so for the limit to pay, the damage of an idle well-year from all sources together, methane, groundwater, property values and orphaning, must be about 37 times its methane damage in oil and 57 times in gas. They follow from Table 9 by dividing the deadweight loss (\$ million) by the idle well-years avoided (thousands), which gives thousands of dollars; Table 12 lists D^* directly for other levels.

A small price instrument always passes. The deadweight loss of a tax grows with the square of its level, like a Harberger triangle, while the abatement grows with the level, so D^* is about half the tax: a tax of \$0.25k per idle year has $D^* =$ \$0.12k in oil. Hence a tax of \$0.25k raises welfare if the damage is above \$0.12k, and it is the best tax when the damage is exactly \$0.25k, since the first-best fee equals the damage. Quantity instruments have no small dose: even the loosest limit forces every well idle for that long to decide,

and D^* runs from \$8k to \$10k across the twelve limits in oil and from \$24k to \$40k across the mandate thresholds (gas: \$10k to \$15k and \$48k to \$58k).

Table 12: Break-even damage per idle well-year, by instrument and level

Instrument	Level	Idle well-years avoided (k)	D^* (k\$/idle yr) $\chi = 0$	$\chi = 1$
Oil				
Idling limit	$\bar{\kappa} = 1$ year	219.2	9.6	10.8
	$\bar{\kappa} = 2$ years	210.1	9.2	10.3
	$\bar{\kappa} = 3$ years	199.4	8.9	9.9
	$\bar{\kappa} = 5$ years	176.4	8.5	9.2
Capacity mandate	0.7 BBL/month	89.7	25.6	12.4
	3.3 BBL/month	129.2	26.9	13.1
	10.2 BBL/month	176.2	27.6	13.9
Idling tax	\$0.25k per idle year	6.8	0.12	0.15
	\$1k per idle year	26.2	0.49	0.57
	\$5k per idle year	105.5	2.2	2.6
Plugging subsidy	\$10k per plug	22.8	0.59	0.50
	\$50k per plug	96.5	3.1	2.4
	\$200k per plug	191.5	13.4	7.8
Bond	\$0.25k per well-year	5.8	19.3	16.5
	\$1k per well-year	22.8	19.2	16.2
Gas				
Idling limit	$\bar{\kappa} = 1$ year	148.0	15.1	17.9
	$\bar{\kappa} = 2$ years	139.4	14.3	16.7
	$\bar{\kappa} = 3$ years	130.0	13.7	15.7
	$\bar{\kappa} = 5$ years	110.7	12.6	14.1
Capacity mandate	78.2 MMBtu/month	61.8	47.6	20.2
	195.8 MMBtu/month	91.9	48.4	21.7
	403.1 MMBtu/month	122.0	50.3	23.3
Idling tax	\$0.25k per idle year	3.0	0.12	0.15
	\$1k per idle year	11.6	0.49	0.60
	\$5k per idle year	50.4	2.3	2.8
Plugging subsidy	\$10k per plug	9.6	0.63	0.50
	\$50k per plug	43.4	3.3	2.4
	\$200k per plug	108.9	14.4	8.7
Bond	\$0.25k per well-year	2.4	33.4	26.7
	\$1k per well-year	9.6	33.3	26.5

Net 2019 dollars; mandate levels are capacity thresholds.

Figure 11 shows the same numbers for every level of every instrument. It is read like Figure 10, with the break-even damage D^* in place of the deadweight loss and with both axes on a log scale: each point is one policy, the horizontal position is its abatement, and moving right along a curve raises the level of the tax, the subsidy or the bond. On the dotted line at the abatement of the limit of 2 years, the limit sits at \$9.2k in oil, the tax that achieves the same abatement at \$7.2k and the subsidy at \$19.7k. The dashed horizontal line is at \$0.25k, the methane damage of an idle well-year: a policy whose point lies below it pays for itself on methane alone, and a reader with a larger damage estimate can move the line up. The left panels count only idle years ($\chi = 0$); the right panels count producing years equally ($\chi = 1$), so their horizontal axis is $\Delta\mathcal{B}_1 = \Delta\mathcal{I} + \Delta\mathcal{Y}$ and the fee and the subsidy lie on the same curve. Three patterns appear. First, price

instruments start near zero at small levels and rise with the level. Second, the quantity instruments sit well above the methane line at every level. Third, when producing years count, the subsidy and the bond, which induce plugs, become cheaper relative to the limit, which returns wells to production.

The ordering is the same when we re-estimate oil on single-well leases (Section 4.4): at the abatement of the limit of 2 years, the limit costs 1.32 times the fee and the subsidy 2.02 times, against 1.29 and 2.75 in the pooled oil fit.

5.5 What a given damage justifies

We now fix the damage of an idle well-year, whatever its source. Table 13 evaluates the headline instruments at damages per idle well-year from \$0.10k to \$2k, at $\chi = 0$. Methane alone gives the first four values, which correspond to leak rates of 6.2, 10, 15 and 27 g/h; the fourth, 27 g/h, is the mean leak of orphaned wells in Louisiana. The last two values, \$1k and \$2k, illustrate damages that add other channels to methane. For each damage, the first column gives the idle well-years avoided by the first-best tax, $\tau_s = d_s$, and the other columns give the net benefit $d_s\,\Delta\mathcal{I} - \text{DWL}$ of that tax and of the other instruments at their headline levels: a \$20k subsidy, the limit of 2 years and the mandate at 3.3 barrels a month. At the methane damage of \$0.25k, the first-best fee of \$0.25k per idle year avoids 6.8 thousand discounted idle well-years in oil, 3.0% of the baseline (3.0 thousand and 1.9% in gas), and its net benefit is \$0.8 million in oil and \$0.4 million in gas. At the same damage, the net benefit of the limit of 2 years is −\$1.9 billion in oil, that of the mandate −\$3.4 billion, and that of a subsidy of \$20k per plug −\$42 million. A smaller subsidy can pass, as any small enough price instrument does: in oil, only the smallest subsidy on our grid, \$2.5k per plug, has a positive net benefit. We use \$20k as the headline subsidy because it is close to what the RRC spends to plug a well, \$17,012. Even at an idle damage of \$2k, eight times the methane value, the limit's net benefit is −\$1.5 billion in oil and the mandate's −\$3.2 billion. At methane damages, then, the quantity instruments destroy surplus on the order of a billion dollars to avoid damages worth tens of millions, and their case has to rest on damages the methane figures omit, such as groundwater contamination, property values and the fiscal cost of orphaned wells. We do not have estimates of these.

The bond (Table 14) behaves as a plugging subsidy, but its whole payment is a loss, since the premium or the forgone interest goes to no one in this market. Its break-even damage is therefore high: \$19.3k per idle well-year in oil and \$33.3k in gas at $\chi = 0$, about 77 and 133 times the methane damage, nearly the same for every τ_b between \$0.25k and \$1k. When producing years count equally, it is \$16.4k in oil and \$26.6k in gas.

Figure 11: Break-even damage D^* per idle well-year against abatement $\Delta\mathcal{B}_\chi$, by instrument and level: oil in the top row, gas in the bottom row; left, only idle years do damage ($\chi = 0$); right, idle and producing years do equal damage ($\chi = 1$). Curves trace the tax, the subsidy and the bond as their level rises from left to right; points are the idling limits and the mandate thresholds. The dotted vertical line is at the abatement of the limit of 2 years, with the matching tax and subsidy marked. The dashed line is \$0.25k, the methane damage of an idle well-year

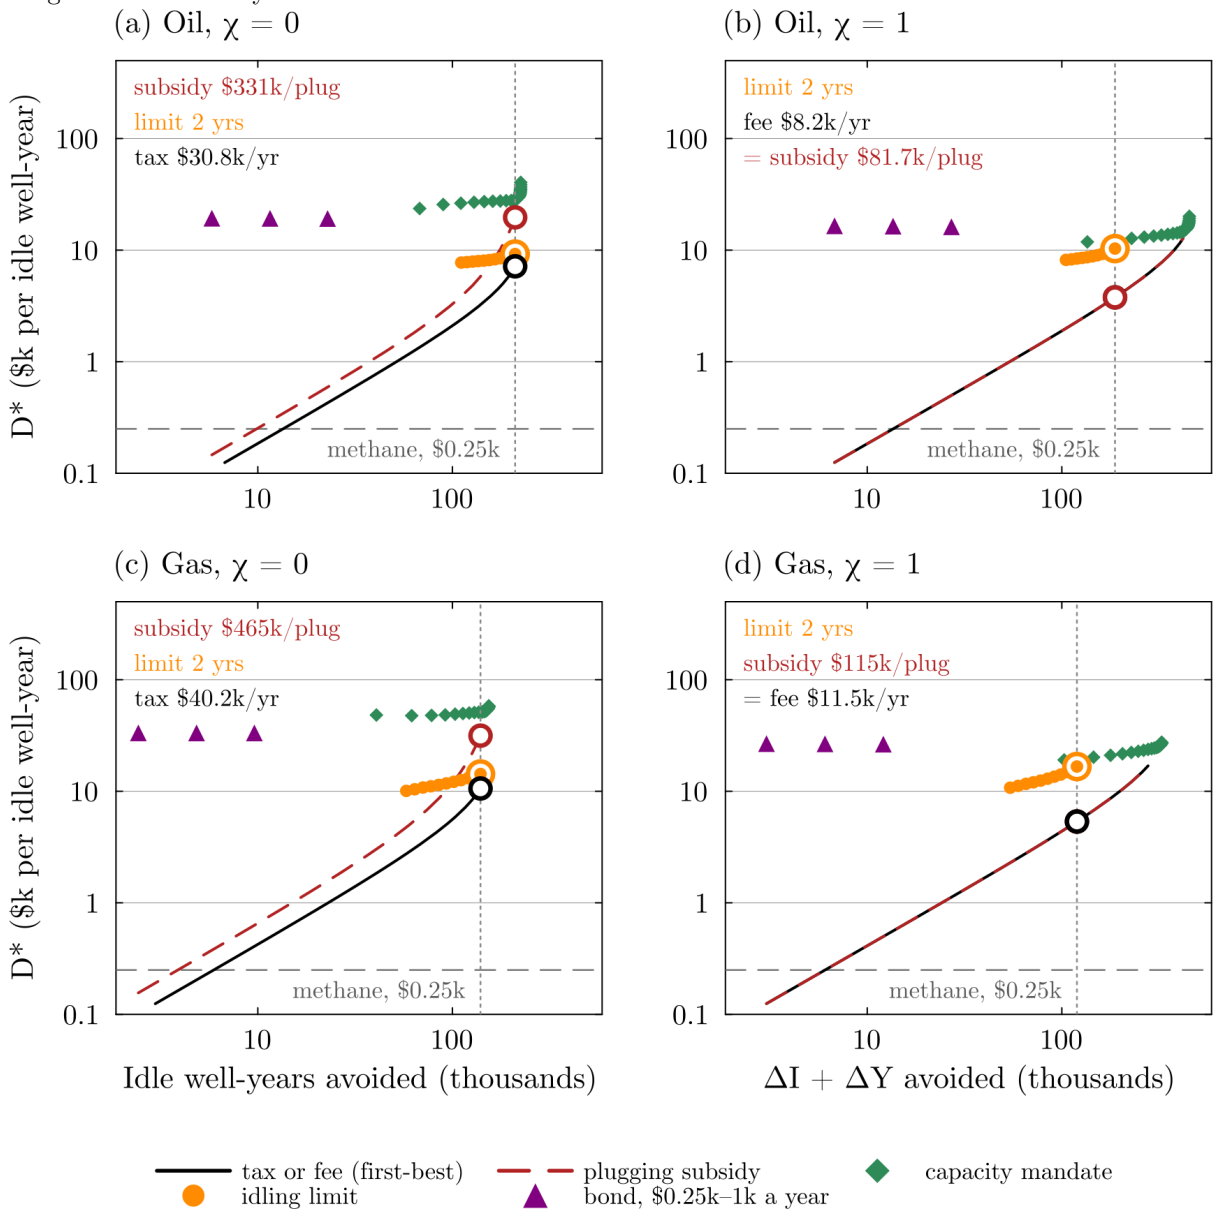

Table 13: Net social benefit at given damages per idle well-year ($\chi = 0$)

Damage d_s	Idle well-years	Net social benefit ($ million)			
(k$/idle yr)	avoided by the tax (k)	tax $\tau_s = d_s$	subsidy $20k	limit $\bar{\kappa} = 2$	mandate
Oil					
0.10	2.7	0.1	-48.7	-1,916.2	-3,455.0
0.16	4.3	0.3	-46.1	-1,903.6	-3,447.3
0.25	6.8	0.8	-42.1	-1,884.7	-3,435.7
0.44	11.8	2.6	-33.7	-1,844.8	-3,411.1
1.00	26.2	13.3	-8.9	-1,727.1	-3,338.8
2.00	50.1	51.7	35.4	-1,517.1	-3,209.6
Gas					
0.10	1.2	0.1	-22.0	-1,977.6	-4,433.5
0.16	1.9	0.2	-20.9	-1,969.2	-4,428.0
0.25	3.0	0.4	-19.2	-1,956.6	-4,419.7
0.44	5.2	1.1	-15.6	-1,930.2	-4,402.3
1.00	11.6	5.9	-5.1	-1,852.1	-4,350.8
2.00	22.5	23.0	13.7	-1,712.6	-4,259.0

Net benefit is $d_s \Delta \mathcal{I} -$ DWL, net 2019 dollars.

Table 14: The bond as a yearly compliance cost on every unplugged well

Yearly cost τ_b	Idle well-years avoided	Break-even damage D^*	
(k$ per well-year)	(thousands)	$\chi = 0$	$\chi = 1$
Oil			
0.25	5.8	19.3	16.5
0.50	11.6	19.3	16.4
1.00	22.8	19.2	16.2
Gas			
0.25	2.4	33.4	26.7
0.50	4.8	33.3	26.6
1.00	9.6	33.3	26.5

Net 2019 dollars. D^* is in k$ per idle well-year; the whole cost of the bond counts as a loss.

5.6 A planner who ignores the heterogeneity

Consider two planners who set the same instrument to reach the same target (Table 15). The smart planner uses the estimated model, with its types and failure state. The naïve planner uses the homogeneous model estimated on the same wells (column (1) of Table 8), in which every idle well is the average idle well. Each planner sets the level of its instrument so that, in its own model, the target is met exactly. We then evaluate both choices in the estimated model, which we take as the truth. The targets are reductions in discounted idle well-years relative to no policy: the reduction the limit of 2 years achieves, and reductions of 10%, 25% and 50% of the baseline idle well-years.

For a target of a 10% reduction in oil, the smart planner sets an idling tax of \$0.8k per idle year and incurs a deadweight loss of \$9 million. The naïve planner sets \$3.5k per idle year, overshoots the target by 57.8 thousand idle well-years and incurs \$130 million. With the subsidy, the smart planner sets \$9.8k per plug and the naïve planner \$58.7k, and the loss is \$13 million against \$401 million. Across targets and fuels, the naïve level is 2.6 to 6.7 times the smart one and the naïve loss 1.7 to 31 times the smart loss. For the largest target in oil, the limit's reduction of 94% of the baseline, the naïve planner cannot reach the target in its own model at any level we allow, so Table 15 marks those rows and the ranges above leave them out. Because the naïve planner overshoots, its larger loss also buys more abatement; Table 15 reports the overshoot so that a reader can value it at their own d_s.

The homogeneous model gets two things wrong at once. Its single plug probability averages over responsive types and over failed or costly types that do not respond, so it misreads how the stock reacts to an instrument. In addition, it converts dollars into utils at a different rate: in the homogeneous model one util is worth \$95.8k in oil and \$98.4k in gas, against \$24.0k and \$38.5k in the estimated model, so the same dollar tax is a smaller tax in utils, and the naïve planner needs a larger dollar tax to move its model. Both push the naïve planner toward a large instrument, and the responsive wells then overreact. We have not separated the two channels. The smart planner needs no well's type, only the estimated distribution.

6 Conclusion

We recover the plugging costs of operators of low-producing wells in Texas from the timing of their decisions to produce, idle and plug. Our estimates suggest that operators who plug a well pay about \$84k in oil and \$60k in gas, net of royalties and severance taxes. This cost covers whatever the operator bears when the well is plugged, including the restoration of the site, so it is comparable with published costs that include land reclamation, to which it is close, and not with the state's spending on the orphan wells it plugs, which covers plugging alone. These costs, and how they differ across wells, are what a regulator needs to choose among

Table 15: A planner who ignores the heterogeneity

Reduction in idle well-years	Instrument	Level (k$) smart	Level (k$) naïve	Overshoot (k well-years)	DWL ($ million) smart	DWL ($ million) naïve
Oil						
as the limit $\bar{\kappa} = 2$	idling tax	30.8	2,000.0‡	14.3	1,504.6	2,234.1
as the limit $\bar{\kappa} = 2$	plugging subsidy	330.8	5,000.0‡	14.3	4,131.8	8,457.4
10%	idling tax	0.8	3.5	57.8	9.4	130.2
10%	plugging subsidy	9.8	58.7	85.9	13.0	400.9
25%	idling tax	2.3	9.6	97.8	61.4	576.8
25%	plugging subsidy	26.0	153.9	121.5	88.1	1,848.4
50%	idling tax	5.5	24.7	91.4	271.8	1,325.7
50%	plugging subsidy	61.8	380.5	101.3	439.2	4,569.6
Gas						
as the limit $\bar{\kappa} = 2$	idling tax	40.2	241.7	14.3	1,477.4	2,479.0
as the limit $\bar{\kappa} = 2$	plugging subsidy	464.7	3,129.7	14.4	4,393.6	8,799.9
10%	idling tax	1.3	3.5	21.9	10.1	62.5
10%	plugging subsidy	16.2	59.8	35.0	15.8	199.0
25%	idling tax	3.6	9.7	43.0	66.6	340.7
25%	plugging subsidy	43.5	160.4	59.7	108.8	1,116.7
50%	idling tax	8.9	25.1	48.5	299.0	1,036.7
50%	plugging subsidy	105.7	406.1	59.0	557.3	3,855.6

Targets are reductions in discounted idle well-years from the no-policy baseline. Tax levels are per idle year, subsidy levels per plug. Overshoot: idle well-years avoided beyond the target when the naïve choice is evaluated in the estimated model. Net 2019 dollars. ‡ The planner cannot reach the target in its own model below the cap of its search ($2,000k per idle year, $5,000k per plug); the level shown is the cap.

instruments, and the costs published from orphan wells cannot provide them.

Three lessons follow for policy. First, when the damage of a well falls on its idle years, a fee on those years is the cheapest way to shorten them. A tight idling limit, such as the Texas twelve-month rule enforced without its extensions, achieves a large reduction at 1.29 times the fee's deadweight loss in oil, because it cannot tell the wells worth keeping from those that are not; a plugging subsidy costs 2.75 times as much, because it also plugs wells that would have returned to production. Second, a limit cannot be dialed down: the limit of 2 years removes 94% of the discounted idle well-years in oil, and none of the twelve limits we compute pays for itself unless an idle well-year does at least $8k of damage, many times its methane damage. At methane damages alone, only small price instruments pass a benefit-cost test; the case for a binding limit, a capacity mandate or a costly bond rests on damages that are not yet measured, such as groundwater contamination, property values and the fiscal cost of orphaned wells. Measuring those damages is the most useful next step for this policy debate. Third, a regulator who ignores how wells differ sets prices too high and overshoots its target, so the heterogeneity matters for the level of an instrument as well as for the choice among them. Extending the analysis to drilling decisions, to operators' portfolios and bankruptcy, and therefore to the creation of orphan wells, is left for future work.

References

Anderson, S. T., Kellogg, R., and Salant, S. W. (2018). Hotelling under pressure. *Journal of Political Economy*, 126(3).

Arcidiacono, P. and Miller, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867.

Boomhower, J. (2019). Drilling like there's no tomorrow: Bankruptcy, insurance, and environmental risk. *American Economic Review*, 109(2):391–426.

Boothroyd, I. M., Almond, S., Qassim, S. M., Worrall, F., and Davies, R. J. (2016). Fugitive emissions of methane from abandoned, decommissioned oil and gas wells. *Science of the Total Environment*, 547:461–469.

Boyd, J. (2001). Financial responsibility for environmental obligations: Are bonding and assurance rules fulfilling their promise? Working Paper 01-42, Resources for the Future. 71 Pages Posted: 14 Oct 2001.

Brown, J. P., Fitzgerald, T., and Weber, J. G. (2016). Capturing rents from natural resource abundance: Private royalties from U.S. onshore oil & gas production. *Resource and Energy Economics*, 46:23–38.

Caputo, M. R. and Wilen, J. E. (1995). Optimal cleanup of hazardous wastes. *International Economic Review*, 36(1):217–243.

Cleveland, C. (2023). Four million wells and counting: The history of oil and gas drilling in the u.s. `https://visualizingenergy.org/four-million-wells-and-counting-the-history-of-oil-and-gas-drilling-in-the-u-s/`. Visualizing Energy, Institute for Global Sustainability, Boston University. Accessed: 2026-05-29.

Costanza, R. and Perrings, C. (1990). A flexible assurance bonding system for improved environmental management. *Ecological Economics*, 2(1):57–75.

Davis, L. W. (2015). Policy monitor—bonding requirements for u.s. natural gas producers. *Review of Environmental Economics and Policy*, 9(1):128–144.

Deighton, J. A., Townsend-Small, A., Sturmer, S. J., Hoschouer, J., and Heldman, L. (2020). Measurements show that marginal wells are a disproportionate source of methane relative to production. *Journal of the Air & Waste Management Association*, 70(10):1030–1042.

Enverus (2026). Enverus. `https://www.enverus.com/`. Online dataset.

Gerard, D. (2000). The law and economics of reclamation bonds. *Resources Policy*, 26(4):189–197.

Guo, H., Nguyen, H., Vu, D.-A., and Bui, X.-N. (2021). Forecasting mining capital cost for open-pit mining projects based on artificial neural network approach. *Resources Policy*, 74:101474.

Heckman, J. and Singer, B. (1984). A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica*, 52(2):271–320.

Ho, J. S., Shih, J.-S., Muehlenbachs, L. A., Munnings, C., and Krupnick, A. J. (2018). Managing environmental liability: An evaluation of bonding requirements for oil and gas wells in the united states. *Environmental Science & Technology*, 52(7):3908–3916.

Honoré, B. E. (1993). Identification results for duration models with multiple spells. *The Review of Economic Studies*, 60(1):241–246.

Hotz, V. J. and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529.

Hu, Y. and Shum, M. (2012). Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics*, 171(1):32–44.

Kahn, J., Franceschi, D., Curi, A., and Vale, E. (2001). Economic and financial aspects of mine closure. *Natural Resources Forum*, 25:265–274.

Kang, M., Kanno, C. M., Reid, M. C., and Onstott, T. C. (2014). Direct measurements of methane emissions from abandoned oil and gas wells in pennsylvania. *Proceedings of the National Academy of Sciences*, 111(51):18173–18177.

Kasahara, H. and Shimotsu, K. (2009). Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica*, 77(1):135–175.

Keane, M. P. and Wolpin, K. I. (1994). The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte carlo evidence. *The Review of Economics and Statistics*, 76(4):648–672.

Keane, M. P. and Wolpin, K. I. (1997). The career decisions of young men. *Journal of Political Economy*, 105(3):473–522.

Lappi, P. (2018). Optimal clean-up of polluted sites. *Resource and Energy Economics*, 54:53–68.

Lappi, P. (2020). A model of optimal extraction and site reclamation. *Resource and Energy Economics*, 59:101126.

Lebel, E. D., Lu, H. S., Vielstädte, L., Kang, M., Banner, P., Fischer, M. L., and Jackson, R. B. (2020). Methane emissions from abandoned oil and gas wells in california. *Environmental Science & Technology*, 54(22):14617–14626.

Muehlenbachs, L. (2015). A dynamic model of cleanup: Estimating sunk costs in oil and gas production. *International Economic Review*, 56(1):155–185.

National Conference of State Legislatures (2022). State oil and gas bonding requirements. `https://www.ncsl.org/energy/state-oil-and-gas-bonding-requirements`.

Newell, R. G. and Pizer, W. A. (2003). Regulating stock externalities under uncertainty. *Journal of Environmental Economics and Management*, 45(2):416–432.

Railroad Commission of Texas (2025a). RRC full wellbore. `https://mft.rrc.texas.gov/link/650649b7-e019-4d77-a8e0-d118d6455381`. Online dataset.

Railroad Commission of Texas (2025b). RRC wellbore query. `https://mft.rrc.texas.gov/link/b070ce28-5c58-4fe2-9eb7-8b70befb7af9`. Online dataset.

Railroad Commission of Texas (2026a). Plugging cost estimates. `https://www.rrc.texas.gov/oil-and-gas/compliance-enforcement/hb-2259hb-3134-inactive-well-requirements/cost-calculation/`. HB 2259/HB 3134 Inactive Well Requirements. Accessed: 2026-05-30.

Railroad Commission of Texas (2026c). Well distribution tables. urlhttps://www.rrc.texas.gov/oil-and-gas/research-and-statistics/well-information/well-distribution-tables/. Accessed: 2026-05-30.

Raimi, D., Krupnick, A. J., Shah, J.-S., and Thompson, A. (2021). Decommissioning orphaned and abandoned oil and gas wells: New estimates and cost drivers. *Environmental Science & Technology*, 55(15). Policy Analysis.

Rice, P. and Smith, V. K. (1977). An econometric model of the petroleum industry. *Journal of Econometrics*, 6(3):263–287.

Riddick, S. N., Mauzerall, D. L., Celia, M. A., Kang, M., Bressler, K., Chu, C., and Gum, C. D. (2019). Measuring methane emissions from abandoned and active oil and gas wells in West Virginia. *Science of The Total Environment*, 651:1849–1856.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033.

Shafiee, S. and Topal, E. (2012). A new approach for estimating total mining costs in surface coal mines. *International Journal of Mining and Mineral Engineering*, 4(2):119–138.

Stavins, R. N. (1996). Correlated uncertainty and policy instrument choice. *Journal of Environmental Economics and Management*, 30(2):218–232.

Sunset Advisory Commission (2016). Sunset advisory commission staff report 2016–2017: Railroad commission of texas. Technical report, Texas Sunset Advisory Commission. Accessed: 2026-05-29.

Tauchen, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181.

Texas Comptroller of Public Accounts (2026a). Crude oil production tax. `https://comptroller.texas.gov/taxes/crude-oil/`. Accessed 2026-11-04.

Texas Comptroller of Public Accounts (2026b). Natural gas production tax. `https://comptroller.texas.gov/taxes/natural-gas/`. Accessed 2026-11-04.

Texas Legislature (2025). S.B. 1150, 89th legislature, regular session: Relating to the plugging of and reporting on inactive wells. `https://capitol.texas.gov/tlodocs/89R/billtext/html/SB01150F.htm`. Signed June 20, 2025; amendments to Natural Resources Code §89.023 effective September 1, 2027. Accessed 2026-11-04.

Townsend-Small, A., Edgar, A., Fernandez, J. M., Jackson, A., and Currit, N. (2024). High rates of hydrogen sulfide emissions measured from marginal oil wells near austin and san antonio, texas. *Environmental Research Communications*, 6(9):091007.

Townsend-Small, A., Ferrara, T. W., Lyon, D. R., Fries, A. E., and Lamb, B. K. (2016). Emissions of coalbed and natural gas methane from abandoned oil and gas wells in the United States. *Geophysical Research Letters*, 43(5):2283–2290.

Townsend-Small, A. and Hoschouer, J. (2021). Direct measurements from shut-in and other abandoned wells in the Permian Basin of Texas indicate some wells are a major source of methane emissions and produced water. *Environmental Research Letters*, 16(5):054081.

Train, K. E. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press, 2 edition.

U.S. Bureau of Economic Analysis (2026). Value added by industry as a percentage of gross domestic product. https://apps.bea.gov/iTable/?reqid=1603&step=2&Categories=GDPxInd&isURI=1. Interactive Access to Industry Economic Accounts Data, Table TVA110. Accessed: 2026-05-29.

U.S. Energy Information Administration (2025a). Crude oil production, annual. https://www.eia.gov/dnav/pet/pet_crd_crpdn_adc_mbbl_a.htm. Accessed 2026-11-04.

U.S. Energy Information Administration (2025b). Natural gas gross withdrawals and production, annual. https://www.eia.gov/dnav/ng/ng_prod_sum_a_EPG0_VGM_mmcf_a.htm. Accessed 2026-11-04.

U.S. Energy Information Administration (2025c). Table 5.2: Crude oil production and crude oil well productivity, 1954–2024. https://www.eia.gov/totalenergy/data/browser/index.php?tbl=T05.02. Monthly Energy Review, Accessed: 2026-05-29.

U.S. Energy Information Administration (2025d). U.s. oil and natural gas wells by production rate. https://www.eia.gov/petroleum/wells/. Released December 8, 2025; Appendix C. Accessed 2026-11-04.

U.S. Energy Information Administration (EIA) (2024). The distribution of u.s. oil and natural gas wells by production rate with data through 2023. Technical report, U.S. Energy Information Administration.

U.S. Environmental Protection Agency (2023). Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances. Technical report, U.S. EPA, National Center for Environmental Economics, Office of Policy. EPA-HQ-OAR-2021-0317; Tables ES.1, 4.1.1 and A.5.1 give SC-CH4 by year and near-term Ramsey rate, 2020 dollars.

Weitzman, M. L. (1974). Prices vs. quantities. *The Review of Economic Studies*, 41(4):477–491.

Yang, P. and Davis, G. A. (2018). Non-renewable resource extraction under financial incentives to reduce and reverse stock pollution. *Journal of Environmental Economics and Management*, 92:282–299.

Yang, P. and Davis, G. A. (2021). Why don't environmental bonds fully cover reclamation costs? *Energy Policy*, 152:112187.

A Conditional choice probability plots

Figure 12: Conditional choice probabilities for each fuel by quantiles of price

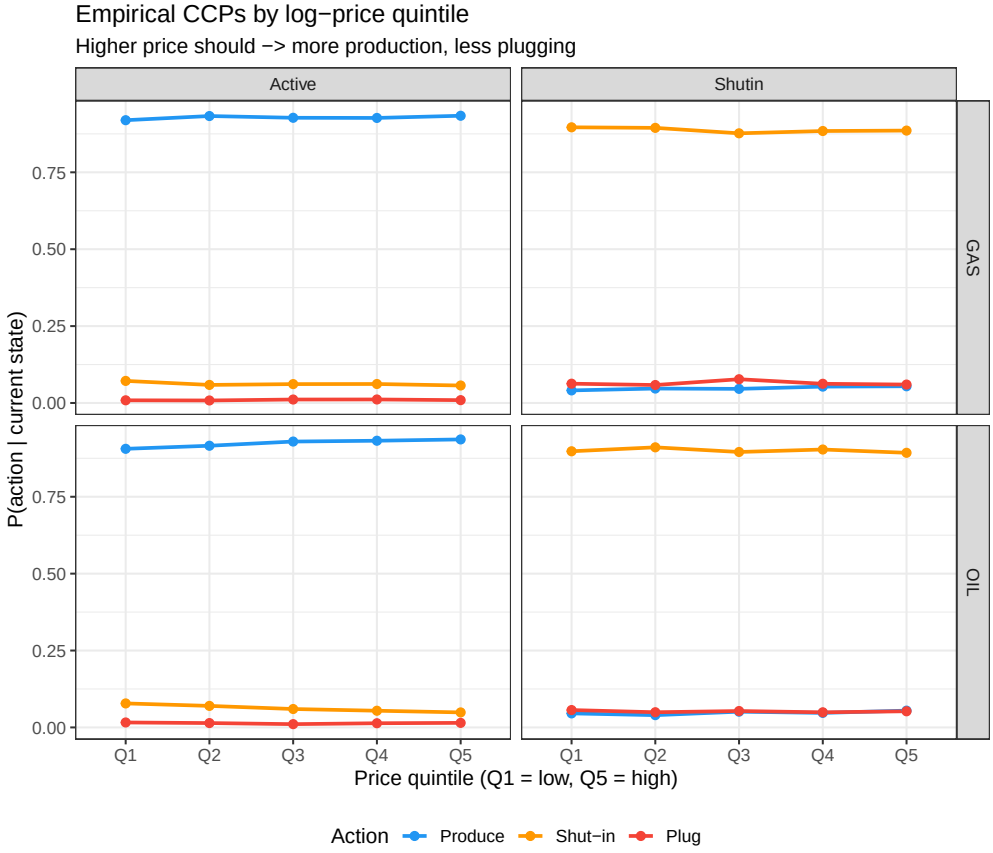

B Methane leak rates

C Identification: the inversion

The argument has two steps. The first recovers, from the panel, the choice probabilities of each type in each failure state, together with the type distribution and the law of motion of the failure state. The second inverts those probabilities, type by type, into the structural parameters, in the manner of Hotz and Miller (1993). We state the first step as an assumption, with the results it rests on, and derive the second exactly. The estimator of Section 3.4 is maximum likelihood, with the first-stage transitions held fixed, and does not follow these steps; they show which features of the data determine each parameter.

Figure 13: Conditional choice probabilities for each fuel by quantiles of depth

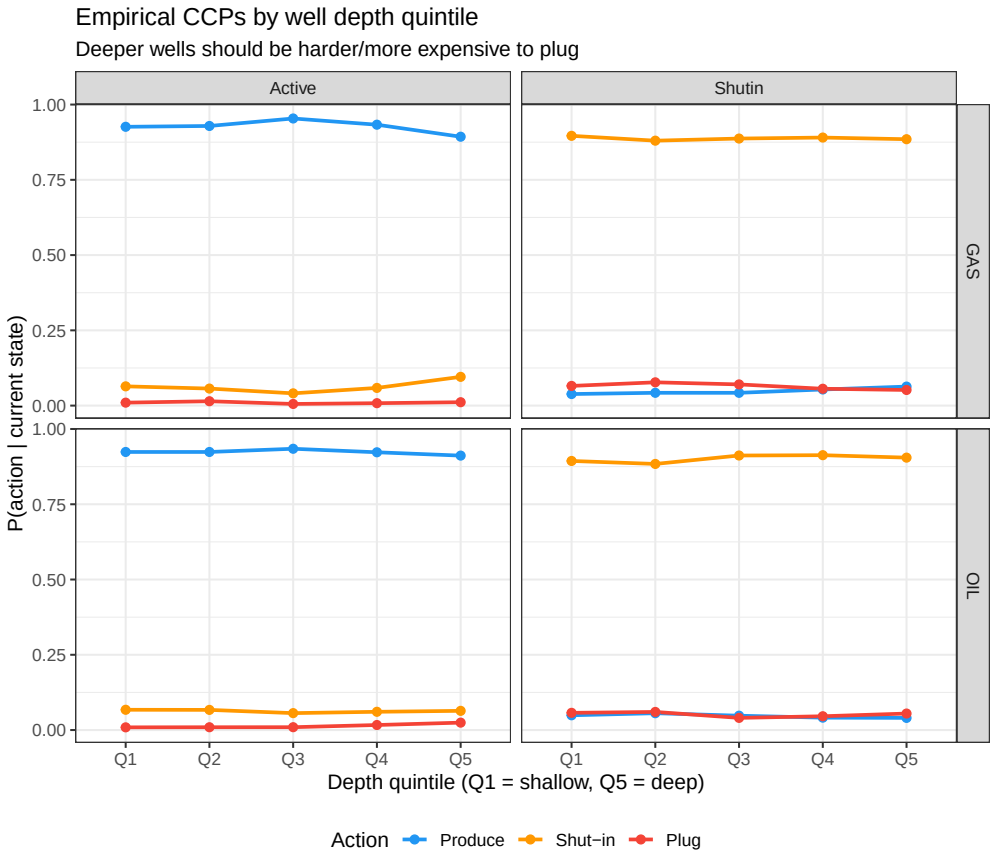

Table 16: Methane leak rates by well status and their damage per well-year

Study	Wells measured	Leak (g CH_4/h)	Damage ($ per well-year)
Idle (unplugged, not producing)			
Townsend-Small and Hoschouer (2021)	Permian, Texas, shut in 3–5 years	6.2	100
Townsend-Small et al. (2016)	WY, CO, UT, OH, unplugged	10.0	162
Lebel et al. (2020)	California, idle	35.4	573
Value we use		15	243
Producing, low production			
Deighton et al. (2020)	Ohio, 0–1 boe a day	128	2,074
Riddick et al. (2019)	West Virginia, active	138	2,236
Lebel et al. (2020)	California, active	189.7	3,073
Townsend-Small et al. (2024)	Luling, Texas, venting	710	11,502

Damage = leak × 8.76 kg a year per g/h × $1,851 per tonne (the EPA 2023 social cost of methane, 2% discount rate, $1,874 in 2020 dollars, U.S. Environmental Protection Agency, 2023, converted to 2019 dollars), that is $16.2 per g/h-year.

C.1 Step 1: choice probabilities by type and failure state

The unobserved state of a well is the pair (q, m): a permanent type q and a binary failure state m that evolves as a controlled Markov chain, frozen after an idle year and drawn afresh with probability $\lambda(o)$ after a producing year. Write $\sigma_a(y, o, w, m; q)$ for the probability of action a at status y and state (o, w) for a well of type q in failure state m.

> **Assumption I.** The type-and-state choice probabilities $\sigma_a(y, o, w, m; q)$, the type distribution, the failure hazard $\lambda(o)$ and the entry distribution of m are identified from the joint distribution of each well's sequence of statuses, actions, capacities and prices.

For a permanent discrete type alone, Assumption I is the finite-mixture result of Kasahara and Shimotsu (2009), which requires a few periods per unit and enough variation in the observed state. For a discrete unobserved state that follows a controlled Markov process, it is the result of Hu and Shum (2012), which requires four or five consecutive periods and rank conditions linking the observables to the unobserved state. The pair (q, m) is a discrete unobserved state of $2Q$ values whose q-component never moves, so it falls under the second result. Although the type is drawn from a LogNormal distribution, the model we estimate is a finite mixture: the likelihood (7) integrates the type over fifteen quadrature nodes with known weights ω_q, so each node is a discrete type, and the LogNormal form ties the costs of the fifteen types together through two parameters, μ_ψ and b_ψ. It is therefore a restricted case of the finite-mixture setting, and the restrictions can only help identification. Two features need care. Latent values are identified only up to their labels, so the types must be ordered, for instance by their reactivation cost, which the LogNormal form does. And plugging ends a well's record, so records have different lengths. We do not test the rank conditions in our data. In the data, the moments that carry this step are the ones Section 3.5 describes: repeated idle spells of the same well for the type, and the weaker price response of wells that stay idle for the failure state.

C.2 Step 2: inverting the choice probabilities

Fix a type q with score z_q and write its costs as $\psi_{ps}(q) = \exp(\mu_\psi + b_\psi z_q)$ and $\psi_{ey}(q) = \psi_{ey}\, e^{\rho_e z_q}$ for $y \in \{p, s\}$. The shocks are Type I extreme value with mean zero, so for every action a,

$$\bar{V}(y, o, w, m; q) = v_a(y, o, w, m; q) - \ln \sigma_a(y, o, w, m; q). \tag{10}$$

Plugging is terminal with $u_e = 0$ and $\bar{V} = 0$ afterwards, so $v_e = -\psi_{ey}(q)$, and (10) with $a = e$ gives

$$\bar{V}(y, o, w, m; q) = -\psi_{ey}(q) - \ln \sigma_e(y, o, w, m; q). \tag{11}$$

The value of a well is minus its plugging cost minus the log of its plug probability, with no Euler constant because the shocks have mean zero.

Substitute (11) into the continuations of (4). After a producing year the well is producing, its capacity and price move by Π_p^{ξ}, and its failure state is 1 with probability $\lambda(o)$; after an idle year it is idle, its capacity and failure state are unchanged, and its price moves by Π^w. Since the rows of the transition matrices sum to one,

$$\begin{aligned} \widetilde{EV}_p(o,w;q) &= -\psi_{ep}(q) - \Lambda_p(o,w;q), \\ EV_s(o,w,m;q) &= -\psi_{es}(q) - \Lambda_s(o,w,m;q), \end{aligned} \tag{12}$$

where

$$\begin{aligned} \Lambda_p(o,w;q) &= \sum_{o',w'} \Pi_p^{\xi}\big((o',w') \mid (o,w)\big)\Big[\big(1-\lambda(o)\big)\ln\sigma_e(p,o',w',0;q) + \lambda(o)\ln\sigma_e(p,o',w',1;q)\Big], \\ \Lambda_s(o,w,m;q) &= \sum_{w'} \Pi^w(w' \mid w)\ \ln\sigma_e(s,o,w',m;q) \end{aligned}$$

are the expected log plug probabilities one year ahead. Both are known after step 1 and the first stage.

Define the adjusted log-odds of producing and of idling against plugging,

$$\begin{aligned} L_p(y,o,w,m;q) &\equiv \ln\frac{\sigma_p(y,o,w,m;q)}{\sigma_e(y,o,w,m;q)} + \beta\,\Lambda_p(o,w;q), \\ L_s(y,o,w,m;q) &\equiv \ln\frac{\sigma_s(y,o,w,m;q)}{\sigma_e(y,o,w,m;q)} + \beta\,\Lambda_s(o,w,m;q). \end{aligned} \tag{13}$$

The log-odds equal $v_a - v_e$. Substituting (4) and (12) gives four equations for each type, each failure state and each (o,w):

$$L_p(p,o,w,m;q) = \theta_R W \times O - \theta_V \cdot O - \psi_m \mathbf{1}\{m=1\} + (1-\beta)\,\psi_{ep}(q), \tag{14}$$

$$L_s(p,o,w,m;q) = -\theta_0 + \psi_{ep}(q) - \beta\,\psi_{es}(q), \tag{15}$$

$$L_p(s,o,w,m;q) = \theta_R W \times O - \theta_V \cdot O - \psi_{ps}(q) - \psi_m \mathbf{1}\{m=1\} + \psi_{es}(q) - \beta\,\psi_{ep}(q), \tag{16}$$

$$L_s(s,o,w,m;q) = -\theta_0 + (1-\beta)\,\psi_{es}(q). \tag{17}$$

The next status is set by the action, not by the origin: a well that produces is a producing well next year, so $\beta\psi_{ep}(q)$ enters both producing equations, (14) and (16), and a well that idles is an idle well next year, so $\beta\psi_{es}(q)$ enters both idling equations, (15) and (17). A well can be producing with $m = 1$: it failed during its last producing year.

C.3 What the equations identify

The left-hand sides are known, so each parameter follows from a difference or a slope of them.

Revenue and production cost. In (14), only $\theta_R \cdot W \cdot O$ and $\theta_V \cdot O$ vary with the state. At a given capacity the slope in the price identifies θ_R, and the remaining variation in capacity identifies θ_V. The same two coefficients must hold in (16), at every type and in both failure states, which over-identifies them. The price varies only over time and is common to all wells, so the slope in w rests on eighteen annual prices; this is why θ_R and θ_V are hard to separate in practice, and why θ_V sits at its bound (Section 4).

Plugging costs. At $m = 0$, what remains of (14) after subtracting the identified revenue and production-cost terms is $(1 - \beta)\psi_{ep}(q)$, which identifies $\psi_{ep}(q)$ since $\beta < 1$. The difference (15) minus (17) is $\psi_{ep}(q) - \psi_{es}(q)$, which then identifies $\psi_{es}(q)$.

Idle cost. Given $\psi_{es}(q)$, (17) identifies θ_0. It must be the same at every type and in both failure states, and (17) must not depend on m, which over-identifies it.

Repair cost. Equation (14) at $m = 1$ minus the same equation at $m = 0$ is $-\psi_m$. A failed and a healthy producing well at the same state have the same continuation after producing, because a repaired well that produces faces the same failure hazard as a healthy one, and the same plugging cost, so they differ only by ψ_m:

$$\psi_m = \ln \frac{\sigma_p(p, o, w, 0; q)}{\sigma_e(p, o, w, 0; q)} - \ln \frac{\sigma_p(p, o, w, 1; q)}{\sigma_e(p, o, w, 1; q)}. \tag{18}$$

Our estimate of ψ_m is very large, so a failed well almost never produces, and (18) then rests on a probability close to zero. The data say that ψ_m is large, not how large.

Reactivation cost and the type distribution. Given the plugging costs, (16) identifies $\psi_{ps}(q)$ at each type. The model says $\ln \psi_{ps}(q) = \mu_\psi + b_\psi z_q$, so two types identify μ_ψ and b_ψ and the other thirteen over-identify them; with the type weights from step 1, this is a test of the LogNormal form.

Type loading. The model says $\ln \psi_{ep}(q) = \ln \psi_{ep} + \rho_e z_q$ and $\ln \psi_{es}(q) = \ln \psi_{es} + \rho_e z_q$, with one ρ_e for both. The slope of the identified $\ln \psi_{ey}(q)$ in z_q identifies ρ_e, and the requirement that both statuses and all types share it over-identifies it.

A rising cost of idling. In a version of the model with an idle clock κ and a step θ_I in the idle cost from year κ_I on, the continuation of an idle year is evaluated at $\kappa^+ = \min(\kappa + 1, \bar{K})$, and (17) becomes $L_s(s, o, w, \kappa, m; q) = -\theta_0 - \theta_I \mathbf{1}\{\kappa^+ \geq \kappa_I\} + (1 - \beta)\psi_{es}(q)$. The step is the jump in this adjusted log-odds at $\kappa^+ = \kappa_I$, within a type. Separating a rising cost of idling from the sorting of types is therefore entirely a matter of step 1: once the choice probabilities of each type are known, the step is read within the type. Section 4.4 reports this specification as a robustness check.

C.4 Null directions

The equations hold with $u_e = 0$ and a zero cost of shutting in a producing well, and with no constant in u_p. A constant would be a fixed production cost.

Claim. Append a constant C_p to u_p. If C_p is fixed, the system (14)–(17) determines $(\theta_0, \psi_{ep}(q), \psi_{es}(q), \psi_{ps}(q), \psi_m)$ uniquely. If C_p is free, it has a one-dimensional null space.

Proof. Hold θ_R and θ_V fixed, since they are identified from slopes, and take a perturbation $(\Delta C_p, \Delta\theta_0, \Delta\psi_{ep}, \Delta\psi_{es}, \Delta\psi_{ps}, \Delta$ at a type q that leaves every left-hand side unchanged:

$$
\begin{aligned}
(14),\ m=0: &\quad \Delta C_p + (1-\beta)\Delta\psi_{ep} = 0,\\
(14),\ m=1: &\quad \Delta C_p - \Delta\psi_m + (1-\beta)\Delta\psi_{ep} = 0,\\
(15): &\quad -\Delta\theta_0 + \Delta\psi_{ep} - \beta\Delta\psi_{es} = 0,\\
(16): &\quad \Delta C_p - \Delta\psi_{ps} - \Delta\psi_m \mathbf{1}\{m=1\} + \Delta\psi_{es} - \beta\Delta\psi_{ep} = 0,\\
(17): &\quad -\Delta\theta_0 + (1-\beta)\Delta\psi_{es} = 0.
\end{aligned}
$$

The first two give $\Delta\psi_m = 0$. The third minus the fifth gives $\Delta\psi_{ep} = \Delta\psi_{es}$.

Case 1, $\Delta C_p = 0$. The first equation gives $\Delta\psi_{ep} = 0$, so $\Delta\psi_{es} = 0$, the fifth gives $\Delta\theta_0 = 0$, and the fourth gives $\Delta\psi_{ps} = 0$. The kernel is trivial.

Case 2, $\Delta C_p = \delta$. The first equation gives $\Delta\psi_{ep} = \Delta\psi_{es} = -\delta/(1-\beta)$, the fifth gives $\Delta\theta_0 = -\delta$, and the fourth gives $\Delta\psi_{ps} = \delta - \delta/(1-\beta) + \beta\delta/(1-\beta) = 0$. The null direction is

$$
(\Delta C_p, \Delta\theta_0, \Delta\psi_{ep}, \Delta\psi_{es}, \Delta\psi_{ps}, \Delta\psi_m) = \delta \cdot \Big(1,\ -1,\ -\tfrac{1}{1-\beta},\ -\tfrac{1}{1-\beta},\ 0,\ 0\Big).
$$

It raises the flow payoff of producing and of idling by δ a year and lowers both plugging costs by the present value $\delta/(1-\beta)$, so every option gains the same present value and no log-odds changes. □

The direction shifts the plugging cost of every type by the same amount. Under the parametric form $\psi_{ey} e^{\rho_e z_q}$ that shift is feasible only when $\rho_e = 0$, so with $\rho_e \neq 0$ the functional form would formally pin down C_p; the direction is null in the type-by-type costs, and we do not rely on the functional form to close it. The estimated model has no fixed production cost, which closes it.

D Notation

Table 17: Notation of the structural model and the counterfactuals

Symbol	Meaning
State and actions	
$a \in \{p, s, e\}$	action: produce, keep idle (shut in), plug (exit)
$y \in \{p, s\}$	status at the start of the year, the action of the year before
o, O	log production capacity and capacity (thousand BBL or MMBtu a year)
w, W	log real price and price (2019 dollars per BBL or MMBtu)
$m \in \{0, 1\}$	hidden failure state, 1 if the well has failed
z_i, q	type score of well i, $N(0, 1)$; quadrature node of the type
Π^o_a, Π^w, Π^ξ_a	transition matrices of capacity, price and both
$\lambda(o)$	yearly failure hazard of a producing well
Payoffs and parameters	
θ_R	utils per thousand dollars of gross revenue
θ_V	cost per unit of capacity
θ_0	yearly cost of an idle well
ψ_{ep}, ψ_{es}	plugging cost from a producing and from an idle well, median type
ψ_{ps}	reactivation cost, $\exp(\mu_\psi + b_\psi z_i)$
ρ_e	loading of the plugging costs on the type score
ψ_m	repair cost of a failed well
π^s_0	share of wells entering idle that have failed
$\check{\psi}_e$	realized plugging cost, conditional on plugging
s_R	share of revenue paid as royalties and severance tax
β	yearly discount factor, 0.90
ε_a, σ_a	taste shock and conditional choice probability of action a
v_a, $\bar{V}$	conditional value of action a and expected value of the well
Policies and welfare	
ς	plugging subsidy, dollars per plug
τ_s, τ_p	fee per idle and per producing well-year
$\bar{O}$	capacity threshold of the mandate
$\bar{\kappa}$	idling limit, consecutive idle years
τ_b	yearly cost of the bond per unplugged well
$\mathcal{I}$, $\mathcal{Y}$	discounted idle and producing well-years of the 2018 stock
$\Delta\mathcal{I}$, $\Delta\mathcal{Y}$	reduction relative to the baseline with no policy
d_s, $d_p = \chi d_s$	damage per idle and per producing well-year
$\Delta\mathcal{B}_\chi$	abatement in idle-year equivalents, $\Delta\mathcal{I} + \chi\Delta\mathcal{Y}$
$\mathcal{U}$, G	operators' value and government revenue
DWL	deadweight loss, the fall in $\mathcal{U} + G$ relative to the baseline
D^*	break-even damage per idle well-year, DWL/$\Delta\mathcal{B}_\chi$