

Relinquishing Riches: Auctions vs Informal Negotiations in Texas Oil and Gas Leasing *

Thomas R. Covert [†]
Richard L. Sweeney [‡]

July 2022

Abstract

This paper compares outcomes from informally negotiated oil and gas leases to those awarded via centralized auction. We focus on Texas, where legislative decisions in the early twentieth century assigned thousands of proximate parcels to different mineral allocation mechanisms. We show that during the fracking boom, which began unexpectedly decades later, auctioned leases generated at least 55 percent larger up-front payments and 40 percent more output than negotiated leases did. These results suggest large potential gains from employing centralized, formal mechanisms in markets that traditionally allocate in an unstructured fashion, including the broader \$3 trillion market for privately owned minerals.

JEL Codes: D44, L13, Q35

Keywords: Auctions, Decentralized Markets, Oil and Gas Exploration and Production

*Both authors declare they have no interests, financial or otherwise, that relate to the research described in this paper, nor do they have any current ties, directly or indirectly to the energy industry. We thank participants at numerous seminars, as well as Steve Cicala, Piotr Dworzak, Aaron Flaaen, Timothy Fitzgerald, Ryan Kellogg, Peter Maniloff, and Cythina Lin Lawell for helpful comments. Yixin Sun, Eric Karsten, Devin McNulty and Grace Park provided excellent research assistance. Code for replication available at https://github.com/rlsweeney/public_cs.texas.

[†]University of Chicago Booth School of Business and NBER, Thomas.Covert@chicagobooth.edu.

[‡]Boston College, sweeneyr@bc.edu.

1 Introduction

Asset owners often need to identify and choose between potential contracting partners to monetize their asset’s value. For example, companies that are the target of acquisition may have multiple potential acquirers, and research institutions looking to commercialize intellectual property often decide among several interested parties. Many land transactions also look like this. How should an owner go about this process? Decades of economic theory have characterized the relative performance of different *formal* transaction mechanisms, like simultaneous bidding in auctions, bargaining following an auction, and pure sequential bargaining (Bulow and Klemperer, 1996, 2009; Roberts and Sweeting, 2013). However, in the real world, many important assets are allocated via *informal*, unstructured processes, and little is known about how they perform relative to theoretical benchmarks.

In this paper, we directly measure the gains from using a centralized, theoretically high-performing mechanism, relative to using informal, decentralized transactions, in the market for mineral leases in Texas. We study a large class of lands initially set aside for public use under the Texas Constitution, on which legislative decisions made nearly one hundred years ago determined whether leases signed during the recent shale boom transacted using an auction or an informal “negotiation.”¹ All of the minerals within these lands belong to the state. On some of these parcels, the state allocates mineral leases using a first price auction. However, on other parcels, the Texas Relinquishment Act of 1919 grants today’s private surface owners the right to negotiate terms with oil and gas companies on behalf of the state, in exchange for half of the revenues they generate. Conversations with many parties involved in Texas leasing confirm that these negotiated leases for public minerals represent a useful analogue to the broader universe of negotiated leases for private minerals in the United States.

Our empirical strategy compares auctioned and negotiated leases that lie in narrowly defined geographic areas, which transact at approximately the same time. Within these location and time bins, the resource quality is the same, the information about its production potential is constant, and, as we argue in Section 4, the allocation mechanism is as good as randomly assigned. Using detailed data from over fifteen hundred leases for publicly owned minerals in Texas between 2004 and 2016, we find that auctioned leases sell for 53 log points more than similar negotiated leases do. This result is robust to a wide range of controls and sample restrictions, and is not caused by differences in the likelihood that auction and negotiation parcels have a successful transaction. The economic significance of these results

¹Throughout the paper we use the term *negotiation* to refer to the informal search, bargaining and solicitation process that lessors use to award drilling rights on private land. We describe what is known about this process in Section 2.

is large: on average, negotiated leases would have generated \$185,000 more up front revenue had they transacted using an auction.

In principle, this improvement in seller revenues could simply reflect a transfer from buyers to sellers. However, one theoretically appealing feature of a well designed auction is that it should allocate an asset to its highest value user. Using the same empirical strategy, we look for evidence of such allocative efficiency gains by comparing subsequent output across auctioned and negotiated leases. We find that auctioned leases produce 39 log points more output than negotiated leases. Combined with the fact that they also have slightly higher royalty rates, we estimate that, on average, auctions increase *total* payments to sellers by about \$341,000 per lease.

Although we are primarily interested in the differences between auctions and negotiations, which are fundamentally transaction level mechanisms, our identifying variation, the Relinquishment Act and privatization policy changes that followed, occurs at the parcel level. If a parcel's lease mechanism affects the likelihood of a successful transaction, then a comparison of auctions vs. negotiations must account for the missing leases that didn't happen in one mechanism but would have under the other. To address this concern, in section 6 we show that negotiation and auction *parcels* are equally likely to be under lease at any point in time. We also replicate our lease regressions in the parcel cross-section, and find that auction parcels earn higher discounted payments and produce more discounted output, inclusive of parcels which never lease.

We then explore the sources of these gains. As we discuss in detail below, there are many sources of persistent differences in productivity across firms in this industry. Nevertheless, we show that while auctions allocate minerals to different firms than negotiations do, both the payment and output differences hold *within* firm, suggesting that firm-lease "match" plays an important role in determining lease outcomes. We corroborate the importance of this channel using data on the full set of bids submitted for auctioned leases. There is a large gap between the winning and losing bids in a lease's auction, of comparable size to the average difference between bonus payments on auctioned and negotiated leases. We also find that there are large gaps in structural estimates of bidder *values* implied by these bids. The difference between winning and losing values is correspondingly similar to the average difference between output on auctioned and negotiated leases. Overall, the distribution of bids and implied values within an auction suggest that firm-lease match is quite important, and that our empirical strategy measures plausible differences in payment and value that a well-designed auction can deliver, relative to negotiations.

These results contribute to a small empirical literature which compares the performance of auctions to alternative mechanisms in real world settings. [Roberts and Sweeting \(2013\)](#)

find that a sequential mechanism performs better than a simultaneous bid auction in timber sales, and [Gentry and Stroup \(2018\)](#) find that the two mechanisms perform similarly in corporate takeovers.² [Salz \(2022\)](#) estimates large returns from the use of intermediaries, who organize auctions on their client’s behalf, in the highly decentralized market for waste collection in New York City. In each of these papers, only one mechanism is observed in the data. To predict what would happen under a different mechanism, the authors estimate the distribution of preferences and costs using a structural model, and then compute counterfactual market outcomes under alternative mechanisms. In contrast, we observe the results of auctions and informal negotiations on otherwise identical objects, transacting at the same time. As a result, we can estimate the causal effects of mechanism choice on welfare relevant outcomes *directly* by comparing observable outcomes of import, like seller revenue and lease output.

A direct comparison between auctions and negotiations is useful because real-world negotiations are messy. Conversations with industry participants suggest that informal mineral lease transactions are heterogenous, with aspects of sequential negotiation coexisting with costly landowner search effort ([Hortaçsu and Syverson, 2004](#); [Allen et al., 2019](#); [Cuesta and Sepulveda, 2018](#)), bilateral bargaining ([Backus et al., 2019, 2020](#); [Larsen, 2021](#)) and even take-it-or-leave-it behavior on the part of some buyers.³ The informality and heterogeneity in this market make it unlikely that the forces driving our results would be captured by any specific formal model of negotiations. Thus, we do not interpret the evidence in this paper as validating or rejecting any existing theoretical comparisons between auctions and specific alternative mechanisms.⁴ Instead, our contribution is to demonstrate the magnitude of the gains from using a formal, implementable, theoretically attractive mechanism, in a real world setting. In that sense our paper is similar to [Larsen \(2021\)](#). Using data from bilaterally bargained vehicle transactions, he finds that the difference in welfare between actual bargaining behavior and “second best” mechanisms achievable under incomplete information is much larger than the [Myerson and Satterthwaite \(1983\)](#) welfare loss arising from incomplete information itself.

We also contribute to the large literature on the economics of oil and gas leasing in the

²There is also a corporate finance literature on mergers and acquisitions comparing auctioned and negotiated outcomes. [Subramanian \(2007\)](#) finds that “Go-shop” deals, in which private equity target firms are explicitly allowed to solicit outside bids following a negotiated acquisition offer, sell at higher prices than “No-shop” deals do. In contrast, [Boone and Mulherin \(2007\)](#) find that auctioned takeover deals transact at roughly the same prices as negotiated deals do.

³In [Appendix A.4](#) we show that lessors with more experience and more land perform better than other negotiators. However the gap between these more sophisticated negotiators and auctions is still substantial.

⁴Indeed, our results are consistent with a variety of theoretical comparisons between simple auctions and more complicated alternatives. We explore these interpretations in [Section 7](#).

United States. Early work by Kenneth Hendricks and Robert Porter on the performance of auctions for mineral leases in the US Gulf of Mexico focused on the empirical relevance of common values and information asymmetries (Hendricks and Porter, 1996, 1988). They showed that US government auctions captured approximately 100% of the *ex ante* surplus in symmetric information environments, but considerably less in asymmetric information environments. Recent work on mineral lease auctions has sought to separately identify affiliation from synergies between neighboring parcels (Kong, 2020), measure bidder uncertainty about competition (Kong, 2021), and evaluate the choice of “security” sold by the winning bidder to the auctioneer (Bhattacharya et al., 2022). An important distinction between this paper and the existing literature is that our results speak to the performance of the *private* leasing market, which constitutes approximately three quarters of all mineral rights in the United States (currently worth over \$3 trillion).

Finally, our paper contributes to a large literature on the costs and benefits of the shale boom. There is now robust evidence of large negative externalities from fracking (Muehlenbachs et al., 2015; Currie et al., 2017). These costs, most of which are local, must be evaluated against the (ideally local) benefits generated by the use of fracking. Previous literature on the landowner benefits of fracking were limited to analyzing partial or indirect measures of landowner revenues, due to the fact that bonus payments are not publicly recorded (Brown et al., 2016; Feyrer et al., 2017; Bartik et al., 2019). The unique institutional features of our setting allow us to directly observe the full set of payments received by mineral owners, including bonus payments and royalty revenue. We find that bonus payments represent two thirds of total landowner revenue earned to-date for the average lease, and by construction, they are the entirety of landowner revenues for the majority of leases that are never drilled. The fact that local payments could be made substantially larger through a more formal allocation process suggests that feasible changes to the way this market operates could meaningfully shift the local cost-benefit ratio of this technology shock.

The rest of the paper proceeds as follows. In Section 2, we describe the mineral leasing process and provide background information on our natural experiment in Texas. Section 3 discusses the data we use and the filtering criteria we apply to it. Section 4 describes our empirical strategy and identification argument. Section 5 describe our main results, which are at the lease level. We study the extensive margin and present parcel level results in Section 6. In Section 7 we discuss possible mechanisms for our results, before concluding in Section 8.

2 Background

In this section, we first describe the contract structure of a mineral lease, the market institutions for leasing, and the dimensions along which some leases may end up being more or less productive (Section 2.1). Then, in Section 2.2, we describe the history of land and mineral ownership in Texas, where a natural experiment motivates our empirical strategy.

2.1 Mineral Exploration and Production in the United States

The US Energy Information Administration estimates that at the end of 2017, oil and gas companies in the United States had proved reserves of 42 billion barrels of oil and 464 trillion cubic feet of natural gas. As of December 31, 2017, these reserves were worth more than \$4.5 trillion.⁵ Although more than three quarters of these deposits lie in land owned by private individuals (Fitzgerald and Rucker, 2016), landowners must partner with oil and gas exploration and production companies (E&P) to transform their reserves into revenue.

Partnerships between landowners and E&P companies are formalized through mineral lease agreements, which are contracts with three key elements: a *primary term* indicating how much time the lessee (the E&P company) has to drill one or more wells; a *royalty rate* providing the lessor (the landowner) with a share of any realized drilling revenues; and an up-front *bonus payment* to secure the right to explore.⁶ Many leases also include *delay rentals*, which are payments the lessee must make to the landowner in the event that drilling has not begun by the start of intermediate milestone events during the primary term. Finally, some leases have additional contractual clauses regarding operations, cleanup and other landowner protections (Vissing, 2017).⁷

Lessees frequently elect not to drill any wells before the conclusion of the primary term, and even when they do, realized drilling does not always result in economically viable quantities of production. As a result, most leases never receive any royalty revenues, so bonus payments are a particularly important aspect of landowner welfare. Despite their conceptual importance in this market, little is known about the distribution of bonus payments because they are usually not recorded in the mineral leases filed in county registries.

Mineral leases are typically initiated by E&P companies, rather than by landowners. An E&P company will conduct background research and decide to acquire drilling rights in a

⁵According to EIA data, oil prices were \$66.73 per barrel (Brent) and natural gas prices were \$3.69 per million BTU (Henry Hub).

⁶This contract structure has important incentive implications, as positive royalty rates provide incentives for lessees to drill later in the contract, and finite primary terms provide incentives for lessees to drill earlier in the contract. See Herrnstadt et al. (2022).

⁷We study these “lease addenda” formally in Appendix E.

particular geographic location. During this acquisition phase, E&P’s often work through intermediaries known as “landmen.” One reason that E&P companies use landmen is that a given firm’s need for new mineral leases may vary over time, and the skills necessary to find landowners, verify their claim to mineral interests, and convince them to lease can be too expensive for an E&P company to consistently maintain in-house. E&P companies can also use landmen to sign leases on their behalf, keeping the E&P company’s identity secret from potential lessors and from competing firms.⁸

In addition to negotiating leasing terms, the decision of *which* firm to partner with is important. During the recent shale boom, there were thousands of E&P companies operating in Texas. Compared to other capital intensive businesses, the oil and gas industry in America is not concentrated, with the 15 largest firms producing about half of all oil produced in Texas in 2020. The largest firm that year, EOG Resources, represented only 7% of total oil production. Given the large number of firms, one might expect a fairly homogenous industry. However, there are many reasons why a landowner might expect different E&P companies to produce substantially more or less output, and royalties, than others. Aside from the obvious size differences, firms differ in the technical sophistication of their engineering choices, their decisions to hire external service contractors to perform drilling and completion services, and their experience developing wells in a specific location.

2.2 Texas Permanent School Fund

Private mineral rights are a uniquely American phenomenon. When individuals outside of the United States purchase surface rights to a piece of land, local or central governments retain ownership and authority over the minerals underground. Because Texas was originally a Spanish colony, early land transactions in Texas followed a similar pattern: when a private individual bought land, the King of Spain retained the mineral rights.

After declaring independence in the mid 19th century, the *Republic* of Texas appropriated millions of acres of unsettled land for public use and the Texas Constitution of 1876 allocated half of this land to benefit public schools. The rules for transactions on the 8 million acres of land, largely in West Texas, contained in this “Permanent School Fund” (PSF), were formalized in 1895. When PSF land was sold to private citizens, Texas, following its colonial tradition, retained the rights to exploit minerals beneath the surface. The surface owner’s remedy for damages resulting from any mineral exploration and development was a payment

⁸When a landman signs a lease on behalf of an E&P company, it must eventually “assign” the lease back to the company before it begins drilling. The data we study in this setting includes all assignments, so we are usually able to determine the true identity of the E&P company that signs a lease, even when it initially uses a landman.

of just \$0.10 per acre, per year.

When oil was discovered in Texas at the turn of the 20th century, surface owners on PSF land argued that this compensation was inadequate.⁹ To stave off “armed rebellion” by these surface owners against state lessees, the Texas legislature passed the Relinquishment Act of 1919 (Shields, 1981). This law, amended and reinterpreted through more than a decade of subsequent litigation, appointed the surface owner as the minerals leasing agent of the state. The arrangement continues to this day, provided that the surface owner’s parcel had been acquired from the PSF by 1931.¹⁰ As a result, when an E&P company wants to develop a Relinquishment Act (RAL) parcel, it must negotiate a lease with the surface owner, as if the surface owner owned the minerals outright. Once the surface owner and the E&P company reach an agreement, they submit the lease to the Texas General Land Office (GLO) for final approval. If approved, the lessee remits half of the bonus and royalty payments to the state. In exchange for acting as the state’s agent, the surface owner gets to keep the other half.¹¹ Although RAL surface owners are free to decide who to transact with and what the terms of that transaction will be, they owe the state a fiduciary duty, and as such, cannot engage in any self-dealing or “off the books” behavior.¹²

Concurrent with the settlement of the Relinquishment Act, Texas passed the Sales Act of 1931. Under this Act, subsequent land sales out of the PSF transferred nearly all mineral rights to the surface owner, under what became known as the “Free Royalty” system. When an E&P company wants to develop a Free Royalty parcel, it negotiates directly with the parcel’s owner, and neither party needs to seek state approval for the lease. While the surface owner keeps all of the bonus and royalty payments that she is able to bargain for, the firm must make an *additional* royalty payment, equal to a 1/16th share of output, to the state. Because the state doesn’t retain any bonus interest, administrative bonus data for Free Royalty leases do not exist.

Texas continued to sell PSF land under the Free Royalty system until the 1970’s, when the OPEC oil embargo and federal oil regulations prompted wide ranging changes to mineral

⁹Although small quantities of oil were observed in Texas prior to that point, recovery in large quantities had proved elusive prior to the massive gusher well at Spindletop in 1901. This well is largely cited as the advent of the oil age in the United States (Yergin, 2008).

¹⁰The Texas Supreme court finalized their interpretation of the Relinquishment Act on June 29, 1931.

¹¹This arrangement does mean that the surface owner experiences *all* of the effort costs associated with negotiating a lease, but enjoys just half of the benefits of a successful lease negotiation. In Section 7 we explore the role that this agency issue could play in explaining this paper’s results.

¹²This is documented in the RAL Leasing Guidelines published by the GLO, available at <https://www.glo.texas.gov/energy-business/oil-gas/mineral-leasing/leasing/forms/Guidelines-for-Leasing-Relinquishment-Act-Lands.pdf>. It is also specified in both the Texas Natural Resources Code and the Texas Administrative Code. Failure to act in the state’s interest can result in invalidated leases, forfeiture of the right to act as the state’s agent in the future, and fines.

management in Texas. Beginning in 1973, the state explicitly retained full mineral interests in all remaining PSF lands. The state continued to sell surface rights to PSF land, but GLO manages mineral leasing in these and other unsold PSF parcels. Unlike leases on RAL parcels, or leases on the broader population of private land, the state awards leases on PSF parcels in which it has retained full mineral rights using an auction. Bidders compete for leases with a fixed primary term and royalty rate, established by GLO before the auction occurs, so the cash bids are analogous to the bonus payment on a negotiated lease.

3 Data

We obtained data from the GLO on the universe of oil and gas leases on PSF lands where the State of Texas retained full or partial mineral rights ([Texas General Land Office, 2020a,b](#)). This includes leases on RAL parcels, land that was sold by the state prior to 1931, and leases on state auction parcels, land either sold after 1973 or land that still belongs to the PSF.¹³ In most analyses, we focus on leases signed between 2004 and 2016, concurrent with the development of several new shale plays. Our initial dataset includes the shape, location, size, effective date, bonus payment, primary term and royalty rate for 4,460 RAL leases and 694 State auction leases. For all leases that eventually result in drilling, we observe monthly payments for gas and oil royalties up through March 2019. We combine the royalty revenue data with output price¹⁴ information to infer monthly oil and gas production for drilled leases. Note that all administrative payment data on RAL leases, both bonuses and royalties, represents *only* the state's 50% share of total payments. As a result, we double all RAL payments in the analyses that follow.

We spatially intersect this lease level dataset with a map of all parcels in the PSF ([P2 Energy Solutions, 2018](#)). We acquired this map from a private contractor hired by GLO to match historical records of PSF parcel information to GIS representations of those parcels. The map identifies which mineral category each PSF parcel belongs to: Relinquishment Act, Free Royalty or minerals fully retained by Texas. We use this map of parcels to measure differences in the likelihood of a successful lease across negotiation (RAL) and auction (non-RAL) parcels in Section 6.1, as well as to measure parcel level outcomes in Section 6.2.

GLO uses first price, sealed bid auctions to allocate its non-RAL leases ([Texas General Land Office, 2017](#)). It conducts two to four centralized auctions per year, each of which includes hundreds of parcels from the PSF and other publicly owned land funds in Texas.

¹³Texas did not retain primary mineral ownership on Free Royalty lands sold between 1931 and 1973. As a result, GLO does not track full payment information on these parcels.


¹⁴We use the Henry Hub price for natural gas royalties and the WTI price for oil royalties ([U.S. Energy Information Administration, 2019e,b](#)).

For every parcel that is nominated by an E&P company, we observe a “bid notice” describing the parcel itself, the date that the auction will be held, and the reserve price. Following the auction, we observe the name of each bidder who bid above the reserve, as well as their bid.

3.1 Sample Selection and Initial Comparisons

We impose a number of restrictions on these data to obtain our final sample. First, we restrict the sample to leases lying on top of a shale formation, as our empirical strategy leverages the unexpected shock to the value of land from the fracking boom which occurred decades after the Relinquishment Act.¹⁵ We then exclude leases smaller than 10 acres or bigger than 1,000 acres, as parcels in the PSF are never more than 1,000 acres, and GLO rarely auctions leases that cover more than one parcel. We also remove leases where the minerals are negotiated by multiple parties, due to inheritance issues. Finally we exclude a small number of leases with missing or inconsistent information. We provide additional information on these restrictions, as well as an associated waterfall table with the number of leases surviving each cut, in Appendix B.1. The resulting dataset of 1,061 negotiated leases and 454 auctioned leases is summarized in Table A.6. Figure 1 demonstrates the distribution of lease types over time.

Figure 1: Sample Leases by Year and Type



In the cross section, auctioned leases are larger, have slightly “more convex” shapes, and are less likely to cover more than one legally defined piece of land, although the differences in

¹⁵We use the EIA’s definition of shale formations in Texas, shown shaded in yellow in Figure A.1 (U.S. Energy Information Administration, 2019c,a,d).

Table 1: Lease Summary Statistics by Type

Variable	Negotiation (N = 1061)				Auction (N = 454)				Difference	p-value
	mean	sd	min	max	mean	sd	min	max		
Land Characteristics										
Acres	0.29	0.26	0.01	0.99	0.36	0.24	0.01	0.73	-0.07	0.00
Shape Quality	0.94	0.15	0.08	1.00	0.96	0.10	0.24	1.00	-0.02	0.00
Multiple Parcels	0.07	0.25	0.00	1.00	0.03	0.16	0.00	1.00	0.04	0.00
Shale Thickness	3.20	1.23	0.15	5.90	3.77	1.21	0.10	5.84	-0.56	0.00
Lease Characteristics										
Bonus	0.94	1.17	0.01	10.42	2.09	2.69	0.02	19.00	-1.15	0.00
Term	3.84	1.16	1.00	7.00	4.76	0.65	3.00	5.00	-0.93	0.00
Royalty Rate	0.24	0.02	0.19	0.25	0.25	0.01	0.20	0.25	-0.01	0.00
Lease Outcomes										
Output	1.87	4.76	0.00	47.07	2.29	4.98	0.00	31.71	-0.42	0.13
Lease Production Revenue	6.58	16.82	0.00	153.49	8.53	20.07	0.00	169.36	-1.95	0.07
Total Seller Revenue	2.69	4.63	0.02	45.02	4.56	5.99	0.03	44.49	-1.86	0.00

Units: acres are reported in thousands; shale thickness is reported in thousands of feet and is not available for leases overlying the parts of the Eagle Ford shale, nor for any leases overlying the Haynesville and Barnett shales (222 negotiation leases and 33 auction leases); bonus, lease revenue, and total seller revenues are all reported in thousands of nominal dollars per acre; output is reported in hundreds of barrels of oil equivalent per acre; term is reported in years. *Definitions:* shape quality is the ratio of the lease’s size to the size of the convex hull containing it; we define a lease as “drilled” if it ever reports a royalty payment; output is the BTU-weighted discounted sum of oil and gas production observed on the lease, through March 2019, lease production revenue is the monetary value of this production at contemporaneous prices, and seller revenue is the sum of bonus payments, the mineral estate’s royalty share of lease revenues, and other fees.

these measures are small. They also generate substantially higher bonus payments (per acre), pay slightly higher royalty rates, and have longer primary terms.¹⁶ Auctions produce more oil and gas and more associated lease revenue. As a result, the difference in total revenues, the sum of bonus payments, royalty income on production, and other fees, is substantially larger than the difference in bonus payments alone. Figure 1 shows that auctions are not consistently prevalent over time, with the auction share as low as 6% (in 2009) and as high as 50% (in 2011). Much of this variation is driven by differences across plays in the share of land that is governed by the Relinquishment Act and changes in which shale plays are most active over time. Appendix Figure A.1 shows that negotiated leases are considerably more prevalent in all shale plays except the Permian Basin in West Texas. The variation in auction share that Figure 1 documents is consistent with this, as the lower auction share years coincide with peak years for other shale plays, like the Eagle Ford. These differences in timing and location underscore the importance of flexibly controlling for these factors in our empirical specifications below.

To construct parcel level outcomes, we match all leases on PSF land to their respective parcels, using the shapefile provided by P2Energy. When leases span multiple parcels, we allocate lease outcomes to parcels proportionately based on parcel acreage. Table 2 presents parcel level summary statistics. We compute the present value of lease outcomes associated with each parcel, discounting back to 2004 at a ten percent discount rate. On average, RAL “negotiation” parcels are more likely to have a lease, and have more leases over time, than state “auction” parcels.¹⁷ Negotiation parcels generate smaller total bonus payments, but have similar average realized output and production revenue.

¹⁶In principal, RAL surface owners are free to negotiate bonuses, royalty rates, and primary terms as they see fit. However, GLO can always reject lease terms they view as insufficiently competitive, and as a result, they effectively require that negotiated royalty rates and primary terms are similar to those they are setting at auction. Moreover, for all leases sold at the same auction, GLO sets nearly uniform royalty rates and primary term lengths.

¹⁷In some circumstances, state owned minerals are negotiated, and RAL minerals are auctioned. Despite this, in this section, we describe state parcels as being “auction” parcels for ease of exposition, and to make the regression tables comparable to Section 5.

Table 2: Parcel Summary Statistics by Type

Variable	Negotiation (N = 2014)				Auction (N = 607)				Difference	p-value
	mean	sd	min	max	mean	sd	min	max		
Land Characteristics										
Acres	0.36	0.27	0.00	1.29	0.19	0.23	0.00	0.92	0.17	0.00
Shape Quality	0.97	0.10	0.02	1.00	0.95	0.12	0.24	1.00	0.01	0.01
Shale Thickness	3.29	1.36	0.07	6.17	3.37	1.47	0.06	5.84	-0.08	0.35
Parcel Outcomes										
Ever Leased	0.72	0.45	0.00	1.00	0.55	0.50	0.00	1.00	0.17	0.00
Number of Leases	1.72	2.86	0.00	61.00	0.84	0.88	0.00	4.00	0.88	0.00
Bonus	0.39	0.57	0.00	15.78	0.69	1.23	0.00	12.71	-0.30	0.00
Output	1.24	2.97	0.00	41.82	1.13	2.67	0.00	17.30	0.11	0.38
Lease Production Revenue	4.11	9.13	0.00	101.96	4.06	9.94	0.00	68.81	0.05	0.91
Total Seller Revenue	1.46	2.36	0.00	26.46	1.80	3.22	0.00	26.41	-0.34	0.02

Units: acres are reported in thousands; shale thickness is reported in thousands of feet and is not available for parcels overlying the parts of the Eagle Ford shale, nor for any parcels overlying the Haynesville and Barnett shales; bonus, lease revenue, and total seller revenues are all reported in thousands of nominal dollars per acre; output is reported in hundreds of barrels of oil equivalent per acre. *Definitions:* shape quality is the ratio of the lease’s size to the size of the convex hull containing it; ever leased is a dummy equal to 1 if the parcel has any active lease during 2004-2016, including leases signed before 2004; number of leases is the number of leases signed on the parcel between 2004-2013; output is the BTU-weighted discounted sum of oil and gas production observed on the parcel, through March 2019, lease production revenue is the monetary value of this production at contemporaneous prices, and seller revenue is the sum of bonus payments, the mineral estate’s royalty share of lease revenues, delay rentals, and miscellaneous fees.

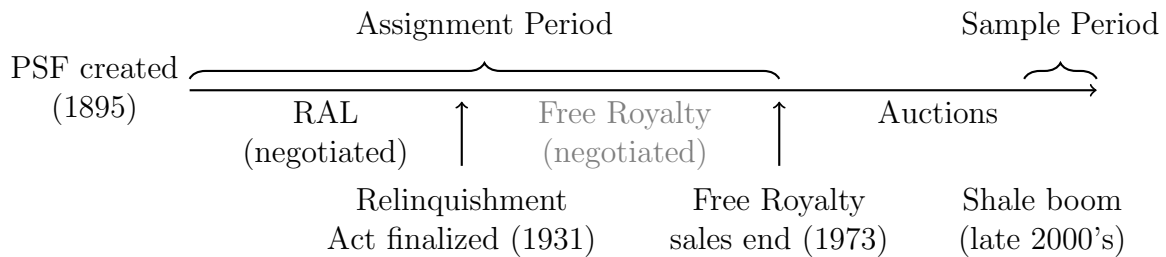
4 Empirical Strategy

The single most important factor determining the value of a lease is the quality of the geological resource underneath it. Leases on parcels located above better mineral resources may transact at higher prices, attract more investment and produce more output. However, we do not observe mineral quality directly, so our empirical strategy compares auctioned and negotiated leases lying in the same narrowly defined geographic area, over which the underlying geological resource is likely to be constant. In an ideal version of these comparisons, we would have randomly assigned mechanism type, auction or informal negotiation, within narrow geographic areas among the population of private mineral owners on top of shale formations, at the eve of the fracking boom. In practice, we observe leases signed between 2004 and 2016 on PSF parcels where the state retained a mineral interest. Within this sample, a lease’s mechanism type is determined not by randomization, but by the date on which the *parcel* underlying it was first sold by the state. Our first identifying assumption is

thus that variation in *parcel privatization* dates, within a narrowly defined geographic area, is uncorrelated with the unobservable geological quality within that area.¹⁸

Figure 2 depicts the timeline of parcel mechanism assignment. All parcels sold out of the PSF prior to 1973 transferred mineral negotiating rights to the buyer: Relinquishment Act lands until 1931, and Free Royalty lands thereafter. However, the state only retained an interest in the minerals underlying Relinquishment Act parcels. Because we observe the full payment structure on present day leases on these lands, but not for leases on parcels sold under the Free Royalty system, leases on Relinquishment Act parcels represent our “control” group. By 1973, the State of Texas ended the practice of selling the rights to minerals when it sold PSF land. Leases on these subsequent land sales, as well as leases on parcels that still belong to the PSF, are awarded using a formal auction, and thus make up our “treatment” group. In light of this history, our identification assumption thus requires that whether or not a parcel was sold before 1973 (so that its leases are negotiated), and, conditional on that, whether it was sold before 1931 (so that we observe leases overlying it), are both uncorrelated with unobserved geological quality, within a narrowly defined geographic area.

Figure 2: Lease Assignment Timeline



One threat to the validity of these assumptions is the possibility that the State of Texas and/or buyers of PSF land had knowledge about which parcels, within narrowly defined geographic areas, would be better or worse for eventual shale development. For example, if land buyers prior to 1973 knew where the “good” parcels were, they might rationally have acquired them quickly, leaving only “bad” parcels for future auctions. Similarly, if the State of Texas had equivalent knowledge and wished to retain “good” parcels for their eventual participation in mineral lease auctions during the shale era, RAL and Free Royalty parcels would be worse, on average. We view both of these scenarios as unlikely. The primary determinants of better vs. worse parcels in the modern shale era are characteristics of the shale rock beneath these parcels, including how thick it is and the concentration of hydrocarbons within it. However, the assignment process is complete by 1973, decades before

¹⁸This empirical strategy is similar to what [Leonard and Parker \(2021\)](#) employ in their study of resource development in the North Dakota Bakken Shale boom.

even the approximate locations of major shale deposits were known or the technology capable of exploiting it was invented. For this reason, it is reasonable to assume that negotiated and auctioned parcels overlie rock that is similarly valuable in the modern shale era. Moreover, since the RAL to Free Royalty transition occurred in 1931, decades before the negotiation to auction transition in 1973, the fact that we can't observe leases on Free Royalty parcels is similarly innocuous.

Though we can't directly test assumptions about the distribution of unobserved quality, we can test whether land sold during the different eras depicted in Figure 2 is similar on observable measures of quality. Table 3 presents a series of balance tests using the entire sample of *parcels* that overlie shale formations in the PSF. We begin by projecting the best available measure of shale resource quality, shale thickness, onto indicators for whether the parcel was sold during the Free Royalty period or whether the minerals were retained by the state, along with location fixed effects. The excluded category is RAL parcels. The small point estimates and precise standard errors in this first regression suggest that within a geographic area, the three types of parcels overlie similarly thick shale rock. This is not surprising, since the locations of thicker vs. thinner parts of shale plays were not known until long after 1973.

Table 3: Parcel comparison: Land in the PSF Overlying Shale Formations

	Thickness	Acres	Shape	Water	Rivers
Auction	0.020 (0.054)	-83.018 (14.906)	-0.008 (0.006)	0.144 (0.368)	0.067 (0.064)
Free Royalty	-0.013 (0.039)	-76.742 (13.451)	0.005 (0.005)	-0.080 (0.327)	0.014 (0.050)
Average	3.069	276.652	0.962	11.772	0.598
N	2,487	3,731	3,731	3,731	3,731
R^2	0.903	0.496	0.344	0.823	0.545

Definitions: Thickness is the thickness of the shale formation in thousands of feet, and is not available for parts of the Eagle Ford shale, nor for any of the Barnett and Haynesville shales. Acres is the size of the parcel, in thousands of acres. Shape Quality is the ratio of parcel size to the size of the convex hull containing the parcel. Water is the distance in thousands of meters from the parcel's centroid to the nearest freshwater lake, pond, marsh or reservoir and Rivers is the distance in thousands meters to the nearest river or stream. All models include fixed effects for the 10 mile grid containing the centroid of the parcel, and standard errors are clustered at the grid level.

We also check whether the three parcel types differ on surface characteristics that are useful for oil and gas development and were known at the time that the mechanism type was determined. Unlike shale rock quality, it is possible that parcels would differ along

surface dimensions, because RAL and Free Royalty purchasers explicitly acquired surface rights for economic use during the pre-shale era. Some surface characteristics might be simultaneously valuable to both pre-shale surface use and shale-boom drilling. For example, bigger parcels or parcels with more convex shapes could be more valuable to agriculture, by making mechanical plowing more efficient, as well as shale development, by making efficient well spacing easier. Similarly, parcels with better access to water could be more valuable in irrigated agriculture, and also more useful in shale development, because water is a key input in the hydraulic fracturing process. The next four models in Table 3 estimate similar regressions, using parcel size, shape and distance to water resources as the outcome variable. Both auction and Free Royalty parcels are smaller than RAL parcels. Because of this, all of our regressions flexibly control for the size of the lease or the parcel. However, all three parcel types have similarly convex shapes and are equally close to water resources (U.S. Geological Survey, 2019).

Having established that a parcel’s assignment mechanism is as good as randomly assigned, conditional on its location, we next discuss the additional assumptions necessary to compare auctioned and negotiated *lease* outcomes. Leases derive part of their value from the mineral resource underneath the parcels they are associated with, so we’ll condition on where leases are in our comparisons between auctioned and negotiated leases. Market conditions are another important component of lease value and performance. Leases signed during periods of high output prices or increased technological progress may be more useful to E&P firms, and may generate better post-leasing outcomes, even after conditioning on location. To account for this, we will also condition on when leases are signed. Many other important determinants of mineral lease values in other settings are essentially fixed among the population of auctioned and negotiated leases on PSF land, for institutional reasons. Leases on both types of parcels are buyer initiated, with an E&P company approaching a landowner or nominating a parcel for auction to GLO. The GLO directly constructs auction reserve prices based on the value of recent nearby sales, and uses the same methodology to reject negotiation offers that they deem to be too low.¹⁹ The GLO also requires RAL surface owners to use a standard lease document, whose structure is nearly identical to the lease contract GLO uses in its auctions, so auctioned and negotiated lease contractual terms are effectively the same.²⁰ Finally, in both auctioned and negotiated leases, the mineral estate belongs to the State of Texas, so

¹⁹Although the auction reserve prices are public, while the negotiation “reserve” prices might not be, the fact that they are both constructed using the same information allays concerns that the reserve may be differentially informative about parcel quality across the two mechanisms.

²⁰This is in stark contrast to the broader private mineral leasing market, where contractual terms vary considerably (Timmins and Vissing, 2022). Bajari et al. (2009) suggest that for complex projects, negotiation may allow flexible contracts that improve ex post value generation. Here, the standardized nature of the lease contracts limits scope for this channel to generate differences across the two mechanisms.

the same government body is the legal representative when conflict arises with lessees.


With these assumptions in mind, we estimate several versions of the following lease-level regression,

$$Y_i = \tau \text{Auction}_i + X_i \beta + \delta_{L(i),T(i)} + \epsilon_i \quad (1)$$

where Y_i is a lease outcome of interest and Auction_i is an indicator that is equal to one if the lease was allocated by auction. X_i includes controls for the lease’s size and, in some specifications, detailed information about how the surface is used, how far the lease is from other potentially valuable features like water and roads, and the quality of the shale rock underlying the lease.

All of our specifications include direct controls for *where* leases are and *when* they transact, which we write as $\delta_{L(i),T(i)}$ in equation 1. We control for a lease’s location using fixed effects for 10 mile by 10 mile square “grids” that contain its centroid. Figure 3 provides a map with hundreds of leases in one of the most concentrated regions of our data, in the southwest portion of the Permian Basin shale play. Dotted lines show the boundaries of the 10 mile by 10 mile grids we use in many of our fixed effect specifications. Most of the grids shown in this map contain both kinds of leases, and in many grids, auctioned and negotiated leases are direct neighbors. We control for a lease’s transaction time using fixed effects for the year-quarter of its transaction date. In some specifications, we restrict comparisons to leases in the same grid and year-quarter.

Figure 3: Example of Sample Lease Type Overlap



There is no *a priori* sense in which a given fixed effect specification correctly controls for the effects of unobserved quality and market conditions on lease outcomes, and, as Figure 3 demonstrates, the grid boundaries we draw are quite arbitrary. We thus estimate several fixed-effects specifications which vary the size of the grids and the extent to which we interact the fixed effects for time and space. We also non-parametrically control for location and time using a novel application of the double/debiased machine learning techniques (DML) developed in Chernozhukov et al. (2018).²¹ Both of these strategies ensure we are making comparisons between leases with similar mineral quality that transact at similar times.

A final concern with interpreting τ in equation 1 as the causal effect of auctioning vs. negotiating a mineral lease is that the data, by construction, is restricted to successful transactions. However, in both mechanisms it is possible that a potential transaction fails, leaving its parcel temporarily unleased. If the probability of a successful transaction varies

²¹Specifically, we estimate a Robinson (1988) style partially linear specification of equation 1. We follow the procedures recommended in Chernozhukov et al. (2018) (equation 3.5, theorem 3.2, definition 3.3, and equation 4.4), finding the value of τ that minimizes the “cross-fitted” empirical analogue of

$$\mathbb{E}[(Y - \gamma(L, T, X) - \tau(D - \delta(L, T, X)))(D - \delta(L, T, X))]$$

where $D = \text{Auction}$, $\gamma(l, t, x) = \mathbb{E}[Y | L = l, T = t, X = x]$, $\delta(l, t, x) = \mathbb{E}[D | L = l, T = t, X = x]$ and we estimate the functions $\gamma(\cdot)$ and $\delta(\cdot)$ using random forests. For more details, see Appendix C.

across mechanisms, then this sort of selection will bias our estimate of τ for the reasons outlined in Heckman (1979). Moreover, when transactions fail, the seller experiences zero revenue, and if failure is differentially likely between auctions and negotiations, then these zeros are a real part of the differences in seller welfare between the two mechanisms.

We address this concern in two ways. First, we ask whether the two mechanisms generate leases at different rates, which is a necessary condition for selection to bias our estimate of τ in equation 1. In Section 6 we show that the two mechanisms have statistically indistinguishable leasing rates for all but two of the fifty-two quarters in our lease sample. In Appendix A.3, we also show the two mechanisms have statistically indistinguishable hazard rates of lease formation. Second, in Section 6 we re-estimate equation 1 at the parcel level. We compute parcel outcomes as the discounted sum of lease outcomes associated with each parcel, among the entire population of PSF parcels. Because this population includes parcels that never lease, these parcel level analyses are immune to the aforementioned selection concern. Moreover, by studying the aggregate discounted outcomes on each parcel, these parcel level analyses directly capture the “zero revenue” extensive margin effects discussed above. As we discuss in Section 6, the parcel level results are qualitatively similar to the lease level results we report in 5.

5 Lease Results

5.1 Bonus Results

We begin by investigating the impact of auctions on up-front payments, estimating several versions of equation 1 with the natural logarithm of bonus payments as the dependent variable. Table 4 presents the results.²² All models include controls for the size of the lease, in acres. In column 1, we include fixed effects for the year-quarter of the lease’s effective date and for the 10 mile square grid containing the lease’s centroid. The interpretation of this estimate is that auctioned leases generate 55 log points more in bonus payments than similar negotiated leases, and this difference is precisely estimated.²³ In column 2, we interact the grid indicators with year of sample indicators, to account for the fact that different locations in Texas were developed at different times in our sample. Even with these interactive fixed effects, the estimated auction coefficient is stable, with auctions paying 53 log points more, and is still precisely estimated. This model, which compares leases for minerals that are

²²In Appendix Table A.1 we report similar results in levels.

²³Note that in percentage terms, this difference is considerably larger than 55%, as $\exp(0.55) - 1 \approx 73\%$. In appendix A.2, we repeat these regressions in levels, with the dollar value of the bonus payments (per acre) as the left-hand side variable.

located at roughly the same place and which transact at roughly the same point in time, is our main specification.

Table 4: Bonus Payments and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Auction	0.55 (0.06)	0.53 (0.06)	0.44 (0.08)	0.52 (0.07)	0.58 (0.05)	0.53 (0.06)	0.58 (0.05)	0.58 (0.10)	0.59 (0.08)
Grid	10	10	10	20	DML	10	DML	10	DML
Time	Q	GY,Q	GYQ	GY,Q	DML	GY,Q	DML	GY,Q	DML
Extra	No	No	No	No	No	Yes	Yes	No	No
Private Only	No	No	No	No	No	No	No	Yes	Yes
N	1,515	1,515	1,515	1,515	1,515	1,260	1,260	1,308	1,308
R^2	0.914	0.961	0.972	0.944		0.954		0.960	

The dependent variable in each regression is the natural logarithm of bonus payments. In columns 1-4 and 6 and 8, the size of the location bins, in miles, are indicated in the “Grid” row, while the structure of the time controls (“Q” for quarter of sample, “GY,Q” for grid-by-year plus quarter of sample, and “GYQ” for grid-by-quarter of sample) are indicated in the “Time” row. Standard errors are clustered by grid in columns 1-4 and 6 and 8. Columns 5, 7, and 9 use a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). For estimation details, see Appendix C. “Extra” controls include shape regularity, a dummy variable for whether the lease spans multiple parcels, surface cover measures, distance to roads and water sources, and shale thickness, which is only available in the Permian Basin and parts of the Eagle Ford shale. Fixed effect models include a spline in lease size and the DML model includes lease size as a random forest covariate. The average negotiated bonus payment is \$945 per acre.

In the next set of columns, we investigate the sensitivity of these results to alternative time-space controls. In column 3, we include location-quarter-of-sample fixed effects to impose more stringent limits on which leases can be compared over time. To ensure that our results are robust to different choices of spatial controls, in column 4 we use 20 mile grids instead of 10 mile grids. In both cases, the resulting estimates are nearly identical to the results in column 2. In column 5, we replace the grid and time fixed effects with a non-parametric control for the lease’s location and time using random forests in a double/debiased machine learning model (DML). Across all of these specifications, we find consistent evidence that bonus payments are substantially larger in auctions than they are in negotiations. Even at the lower end of these estimates, the implications for seller revenue are large. For an RAL lease of average size, switching to an auction would generate a \$185,000 larger up-front bonus payment.

As we discussed in Section 4, our key identifying assumption is that land which was initially owned by the state and sold by 1931 is similarly valuable for today’s hydrocarbon exploration as land from the same allocation that was not sold before 1973. While we

believe it is unlikely that the timing of early land transactions would be correlated with the productivity of shale formations that were unknown until the early 2000's, our empirical specifications include flexible spatial controls to account for any differences in geology across leases governed by the two mechanisms. Moreover, Table 3 shows that the two types of land are also indistinguishable along observable measures of quality that were known at the time the mechanism type for a parcel was determined. However, even if the two parcel types were similar at the time the mechanism type was determined, by construction, they have been exposed to a different history of surface ownership. RAL surface owners have had at least 73 years²⁴ to develop their surface rights, while State auction parcels were privatized no more than 40 years ago, if ever. To the extent that surface investments help or hinder shale development, they would generate differences across parcel types in their value during the shale boom, even if a parcel's leasing mechanism (negotiation or auction) was as good as randomly assigned.

To ensure that such differences do not confound our estimates, we estimate a series of additional specifications that include measures of surface quality related to subsequent surface investment. Using our lease shape files, we compute the quality of the lease's shape as the ratio of its area to the area of the convex hull containing it and separately determine whether the lease spans more than one distinct parcel. Next, we measure the distance of each lease to road infrastructure (Texas Department of Transportation, 2017). Finally, we compute the most common surface coverage characteristics of a lease (U.S. Geological Survey, 2021).²⁵ In the remaining columns of Table 4, we confirm that the bonus results above are robust to including these characteristics as controls. Columns 6 and 7 estimate the grid-year and DML models with additional controls for measures of surface quality, like the convexity of the lease's shape, an indicator for whether the lease spans multiple parcels, the distance from the lease to roads and water infrastructure, and satellite measures of the lease's landcover, as well as the thickness of the shale underlying the lease. Across all of these specifications, we continue to find that auctions pay significantly more than negotiations do.

One remaining potential confounder, which *is* observably different across the two groups of leases, is surface ownership. The Relinquishment Act specifically allows a subset of private surface owners to perform negotiations, so all of our negotiated leases have private surface ownership. In contrast, some auctions occur on PSF parcels that were never sold, and as a result, have state surface ownership. Private surface ownership itself could reduce the value of a negotiated lease if, for example, private surface owners have houses or livestock on their

²⁴Our lease sample starts in 2004, and the latest RAL privatization from the PSF was in 1931.

²⁵Table 3 demonstrates that leases the parcels underlying our sample leases are statistically indistinguishable along these surface measures.

property, or if E&P companies simply face additional constraints on drilling near private citizens, relative to leases where the state controls the surface. If these constraints made negotiated leases more difficult to develop, E&P companies would rationally pay less to lease them, but this difference in payment would not be caused by the difference in mechanisms. This scenario would violate our exclusion restriction, as the difference in bonuses on RAL lands would come from the fact that they have private surface owners, not the manner in which they were allocated.

To ensure that our results are not driven by differences in private surface rights, we restrict our analysis to parcels where the state does not own surface rights. This constitutes leases on land that the State sold to private individuals after 1973, and excludes leases on land that still belong to the PSF. If there are additional costs to developing leases with private surface ownership, we would expect the difference in bonus payments between these leases and leases on RAL parcels to be smaller than the overall difference we observe, when including the full set of auction leases. Columns 8 and 9 of Table 4 present our main bonus regressions estimated on this restricted sample. Despite the fact that these regressions exclude slightly more than half of all auctioned leases, the estimates are still consistent with the results in columns 1 and 5: both point estimates are statistically significant at conventional levels, and they comfortably lie within the confidence intervals for the estimates covering the entire sample. We therefore reject the concern that negotiated leases earn lower bonus payments because they are associated with private surface ownership.

Finally, surface owners of RAL parcels sometimes negotiate additional contractual provisions which deviate from the standard RAL lease, and it could be the case that these additional contractual demands compensate RAL lessors for the lower bonus payments they receive. To test this hypothesis, we collected and digitized data on the auxiliary clauses embedded in each RAL lease. As we document in Appendix E, we find no evidence that variation in the number of additional contractual demands or the relative landowner vs. E&P company “friendliness” of those contractual demands can explain the differences in bonus payments that we observe. Even after conditioning on these additional contractual characteristics, auctioned leases still pay considerably higher bonus payments than negotiated leases do.

Other Lease Terms

As we noted in Section 2.1, all leases are three-part contracts: bonus, royalty, and primary term. While the previous analysis shows that auction bonuses are higher, it could be that RAL leases are compensated for this difference in up-front payments with higher royalties, or

shorter terms.²⁶ In Table 5, we estimate regressions in which royalty and primary term are the left-hand side variable. Across these specifications, auctions generate leases with higher royalty rates, by about 1 percentage point, and longer primary terms, by about 1 year.

Table 5: Royalty Rates, Primary Terms, and Mechanism Type

	(1)	(2)	(3)	(4)	(5)
Auction - Royalty Rate	0.95 (0.13)	0.86 (0.17)	0.90 (0.25)	0.86 (0.14)	0.82 (0.09)
R^2	0.540	0.752	0.846	0.683	
Auction - Primary Term	1.00 (0.10)	0.86 (0.10)	0.85 (0.13)	0.87 (0.09)	0.99 (0.05)
R^2	0.521	0.732	0.835	0.641	
Grid	10	10	10	20	DML
Time	Q	GY,Q	GYQ	GY,Q	DML
N	1,515	1,515	1,515	1,515	1,515

The dependent variable in the top row is the lease’s royalty rate, in percentage points. The dependent variable in the bottom row is the lease’s primary term in years. In columns 1-4, the size of the location bins, in miles, are indicated in the “Grid” row, while the structure of the time controls (“Q” for quarter of sample, “GY,Q” for grid-by-year plus quarter of sample, and “GYQ” for grid-by-quarter of sample) are indicated in the “Time” row. Standard errors are clustered by grid in columns 1-4. Column 5 uses a double/debiased machine learning routine, as recommended in Chernozhukov et al. (2018). For estimation details, see Appendix C. Fixed effect models include a spline in lease size and DML models include lease size as a random forest covariate.

Given these two facts, it is natural to question how much higher bonuses are *conditional* on these other lease characteristics. However, royalty rates and primary terms are determined by the mechanism of interest, so a regression that included them as right hand side variables would not have a causal interpretation.²⁷ To the extent that higher royalties and/or longer primary terms in auctions make lessors better or worse off, we will capture these effects in the next section, where we examine production, royalty revenues, and total seller revenues, in the same causal framework we use here. For completeness, we also report the results of bonus regressions that include a lease’s royalty rate and primary term as right hand side variables in Appendix Table A.2. The coefficients on the auction dummy, while lacking a causal interpretation, are precisely estimated and similar in magnitude to what we report in this section.

²⁶Note that the *equilibrium* effect of higher royalties and/or shorter terms need not actually increase the value that a landowner gets from a lease, due to the moral hazard issues explored in Herrnstadt et al. (2022). Higher royalties and shorter primary terms may reduce the chance the lessee actually drills.

²⁷This is an example of what Angrist and Pischke (2008) term “bad controls.”

5.2 Allocative Efficiency Results

One potential explanation for the large difference in bonus payments between auctions and negotiations is that auctions select more productive firms, and these firms are correspondingly willing to pay more for the option to drill. As we discussed in Section 2.1, there are many sources of productivity differences across E&P companies. Auctions could select better firms by attracting more bidders than negotiations do, or by simply better identifying the best bidder from the same set that participate in negotiations.²⁸ If this were true, we'd expect auctions to not only pay more than negotiations, they ought to produce more, as well. Alternatively, auctions and negotiations might select equally productive firms, but the absence of formality in negotiations and/or greater competition in auctions allows negotiation winners to pay less than auction winners do. Under this hypothesis, we'd expect auctioned and negotiated leases to produce at similar levels, and the difference in bonus payments that we observe would simply reflect a shift in the division of surplus between firms and landowners. To distinguish between these two theories, we measure the differences in realized output between auctioned and negotiated leases.

We begin by looking at differences in production revenues, as recorded in GLO administrative data. Lessees make monthly royalty payments to GLO. We divide these payments by the associated royalty rate to infer the total revenues generated by a lease during a month, then discount the observed stream of monthly revenue back to the lease's effective date. We call this measure *lease revenue*, as it represents the total production revenue that the lease generates. Though our data covers all production through March, 2019, leases are expected to produce output for 20 years or more. This means that the discounted sum of realized revenue is right censored, even for the earliest cohort of leases. To remove some of the bias due to censoring on the extensive margin, we restrict our analyses output and lease revenues to leases whose primary term has concluded by the end of our royalty data, the last of which was signed in 2013. We also include the same temporal controls as in the bonus regressions, to ensure that we are making comparisons between auctioned and negotiated leases that are at least similarly censored.

Table 6 presents the results. Focusing first on panel (a), the first five model specifications are identical to those in Table 4, showing the effects of mechanism type on each output measure under various spatial and temporal controls. The final column is similar to the

²⁸Given that property rights are well defined here, and typically transferable, one might think that any initial misallocation by negotiators would be quickly rectified through secondary markets. However, even absent any transaction costs, asymmetric information between the initial lessee and a potentially more productive alternative lessee will prevent many socially beneficial reassignments. For a theoretical review of informational frictions and the Coase Theorem, see Farrell (1987). Brehm and Lewis (2021) provide compelling empirical evidence on this point in the specific case of mineral leasing.

specification in column 1 of Table ??, and surface characteristic controls. Across these specifications, auctioned leases produce \$2,615 to \$4,558 more in lease revenue, per acre. Though the point estimates are somewhat imprecise in specifications with finer fixed effects or additional covariates, they are consistently large, and represent an economically significant difference in output, relative to the average lease revenues from negotiated leases of \$5,780 per acre.

Although landowners ultimately care about production revenues (the product of quantities and prices), E&P companies have little ability to influence the prices at which these commodities sell. As a result, variation in output may better capture the traditional notion of productivity differences between auction and negotiation winners than variation in revenues does.²⁹ In consideration of this, we divide our product specific production revenue data by contemporaneous oil and gas prices, and define *output* as the sum of these two series, weighing gas production by its energy content in barrels of oil equivalent terms.³⁰ We then estimate the same set of specifications using output per acre as the left hand side variable, which we report in the second row in the top panel of Table 6. Across specifications, auctioned leases produce 82 to 113 additional discounted barrels of oil equivalent, per acre. These estimates are also more precise than the revenue estimates, because they are not affected by variation in prices over time.

The statistical imprecision in some of these estimates may be driven by the fact that the distribution of both output and revenue, even normalized by lease size, is incredibly skewed. For leases that are ever drilled, the difference between the 10th and 90th percentiles of output per acre spans nearly three orders of magnitude. If all leases were drilled, a natural solution to this skewness would be to estimate differences in revenue and output across leases in relative terms, by using the natural logarithm of these terms as the dependent variables. However, fewer than half of leases are ever drilled, and as such generate zero revenue and output in the real sense (i.e., this is not just a selection problem). In this situation, adding a small constant to these zeros to facilitate the logarithmic transformation is unlikely to be innocuous. To make relative comparisons which respect this skewness, and to control for variation across cohorts in the extent of right-censoring, which we expect to have proportionate effects,³¹ we also estimate pseudo-Poisson regression models, which effectively project the logarithm of the expected value of lease revenue (or output) onto mechanism type, covariates, and our

²⁹For a review of this literature, see Syverson (2011).

³⁰We assume 1,000 cubic feet of natural gas production is equivalent to 0.1767 barrels of oil production, per EIA guidelines here: <https://www.eia.gov/tools/faqs/faq.php?id=45&t=8>

³¹Log or Poisson style models will exactly control for censoring under an assumption that lease output declines at a constant proportional rate over time, a common assumption in petroleum engineering called “Arps decline curve analysis.”

Table 6: Lease Revenue, Lease Output, Seller Revenue, and Mechanism Type**(a)** Linear Regression Models

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Lease Revenue	3.47 (1.28)	2.64 (1.43)	2.61 (2.76)	3.60 (1.23)	4.56 (1.60)	2.91 (1.52)
Auction - Output	0.85 (0.31)	0.82 (0.40)	0.92 (0.73)	1.03 (0.30)	1.13 (0.38)	0.93 (0.42)
Auction - Seller Revenue	1.41 (0.36)	1.15 (0.39)	1.30 (0.72)	1.46 (0.35)	1.89 (0.43)	1.22 (0.42)
Grid	10	10	10	20	DML	10
Time	Q	GY,Q	GYQ	GY,Q	DML	GY,Q
Extra	No	No	No	No	No	Yes
N	1,259	1,259	1,259	1,259	1,259	1,012

(b) Pseudo-Poisson Regression Models

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Lease Revenue	0.63 (0.20)	0.39 (0.16)	0.39 (0.24)	0.55 (0.21)	0.48 (0.18)	0.28 (0.16)
Auction - Output	0.64 (0.20)	0.45 (0.17)	0.55 (0.24)	0.62 (0.25)	0.49 (0.17)	0.39 (0.18)
Auction - Seller Revenue	0.57 (0.13)	0.43 (0.12)	0.42 (0.16)	0.53 (0.15)	0.57 (0.12)	0.41 (0.11)
Grid	10	10	10	20	DML	10
Time	Q	GY,Q	GYQ	GY,Q	DML	GY,Q
Extra	No	No	No	No	No	Yes
N	1,183	746	621	980	1,259	589

The dependent variables are the discounted sum of oil and gas production revenue (Lease Revenue), in thousands of dollars, discounted barrels of oil equivalent (Output), in hundreds, and the discounted present value of the sum of bonus payments, delay rentals paid, fees and production royalties (Seller Revenue), in thousands. In the top panel, the estimates come from linear regression models, with outcomes per acre, while in the bottom panel, they come from pseudo-Poisson quasi-maximum likelihood models, with outcomes in levels. In both panels, the estimates use fixed effects in columns 1-4 and 6, where the size of the location bins, in miles, are indicated in the “Grid” row, and the structure of the time controls (“Q” for quarter of sample, “GY,Q” for grid-by-year plus quarter of sample, and “GYQ” for grid-by-quarter of sample) are indicated in the “Time” row. For these models, standard errors are clustered by grid. In both panels, column 5 uses a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). For estimation details, see Appendix C. The fixed effect models include a spline in acres, while the DML models include acres as a covariate in the random forest routine. “Extra” controls include shape regularity, a dummy variable for whether the lease spans multiple parcels, surface cover measures, distance to roads and water sources, and shale thickness, which is only available in the Permian Basin and parts of the Eagle Ford shale. The sample includes all leases whose primary term ends before March, 2019. In the fixed effect models of columns 1-4 and 6 in the bottom panel, the leases from grids and/or time periods with no variation in output are dropped, as the outcome is completely determined. The average negotiated lease generates \$5,780 in lease revenue per acre, and 1,648 of discounted barrels of oil equivalent per acre.

location and time controls:

$$\log \mathbb{E}[Y \mid \text{Auction}_i, X_i, L_i, T_i] = \tau \text{Auction}_i + X_i \beta + \delta_{L_i, T_i}.$$

The bottom panel of Table 6 presents the same specifications estimated using pseudo-Poisson regression models. Because pseudo-Poisson models are not identified within grids and grid-time combinations that have no variation in output, the fixed-effect estimates drop data from some grids and times. As a result, the pseudo-Poisson estimates in different fixed effect specifications are derived from slightly different samples, most of which are smaller than what we report in the top panel. In spite of this, across specifications, we find consistent and mostly precisely estimated evidence that auctioned leases produce more than negotiated leases do. Whether we use lease revenue or output as the outcome, auctions produce 28 to 64 log points more than negotiated leases do. These results do not depend on how we control for location and time, nor whether we include additional covariates.³²

The combined effects of higher bonus payments and more production imply that auctions generate substantially more net benefit for sellers than negotiations do. We measure this in the “Seller Revenue” rows of Table 6, which use the discounted sum of bonus payments, discounted royalty revenues, and other fees as the dependent variable. Auctions generate \$1,146 to \$1,894 more revenue per acre for landowners than the negotiations do. Under our main specification (column 2), this difference is worth about \$341,000 for the average RAL lease. These results show that auctions have an economically enormous impact on sellers.

6 Parcel Results

The results in Section 5 show that auctioned leases pay and produce more than negotiated leases do. However, as discussed in Section 4, it is possible that a parcel’s lease mechanism also affects the probability that a lease is signed at all. If this is the case, then a comparison

³²We showed in Section 5 that auction leases have higher royalty rates than negotiation leases do. Because higher royalty rates reduce firms’ incentives to drill and invest in production, failure to control for royalty rates would cause us to underestimate the impact of auctions on output, all else equal.

However, as discussed above, royalty rates are determined after assignment. To retain a causal interpretation in our regressions, we can only condition on royalty rates if we additionally assume that they are determined independently of the mechanism type. In Appendix Table A.3 we show, under this assumption, these differences in output are even higher when we condition on lease royalty rates and primary terms. If royalty rates *were* determined independently of the mechanism type, this would imply that the results Table 6 are in fact underestimates of the true effect of auctions on output.

In Appendix Tables A.4 and A.5 we also show that these differences in output likely come from both the fact that auction leases are more likely to be drilled *and* the fact that they produce more conditional on drilling. However, we caution that this decomposition does not have a causal interpretation because we do not have an instrument for drilling.

of auctions vs. negotiations must account for the missing leases that didn't happen in one mechanism but would have under the other. In this section, we examine outcomes that occur at the parcel level.

6.1 Extensive Margin

We begin our analysis of parcel outcomes by comparing the incidence of successful lease transactions on auction and negotiation parcels. For auction parcels, we can directly compute the probability of a successful transaction, because we observe the list of parcels available in each auction event, as well as the subsequent bids.³³ However, we don't observe attempted negotiations that fail, so on negotiation parcels, we observe neither the likelihood of nomination nor the probability of a successful transaction after being nominated. In lieu of this, we use the linked parcel-lease data to directly look for differences in the extent to which a parcel is leased under the two mechanisms, inclusive of both differences in the likelihood of nomination and transaction success given nomination.

To analyze this, we project a dummy indicating that a parcel is under lease in a given quarter onto an indicator for auction, as well as size and location controls. Figure 4 provides a visual representation of the results. Though auction parcels are somewhat more likely to be under a lease in most quarters, the individual point estimates are quite noisy. For all but 2 quarters, we cannot reject the hypothesis that auction and negotiation parcels are equally likely to be under a lease. As a result, we conclude that the two mechanisms do not differentially select parcels into leasing, and that the difference in outcomes between auctioned and negotiated leases is not due to a difference in the likelihood of transaction.³⁴


At first glance, the fact that auctions pay more but are no less likely to generate a successful transaction is perhaps surprising. However, this is explainable by two facts that we develop further below. First we show in section 7 that the most active firms in our data participate in both mechanisms. Since there are no institutional restrictions on participation in either mechanism, it seems likely that nearby parcels allocated via different mechanism likely draw from the same distribution of bidders. Second, the same team that sets auction reserve prices at GLO also has final say over whether a negotiated transaction can take place,

³³Among GLO auctions on PSF land, 40% of nominated parcels failed to receive a qualifying bid, so on a per potential transaction basis, failure is quite common. The GLO often offers to sell these failed parcels again in future auctions, to the point that 83% of all observed nominated parcels transact *at some point* in our sample.

³⁴In appendix A.3 we further explore this possibility with a formal duration analysis. Specifically, for each parcel, we identify "unleased spells," or periods during which a parcel has no active lease, and then ask whether RAL and State parcels in similar places, of similar sizes, that begin unleased spells at similar times, end their period of being unleased at similar rates. Our analyses show that we cannot reject the hypothesis that unleased spells on RAL and State parcels end equally quickly.

so it seems likely that the two mechanisms have similar reserve prices. Later on, in Table 12, we show that while negotiated bids are much lower than nearby *winning* auction bids, they are in fact higher than nearby reserve prices and statistically equivalent to the nearby minimum bid placed. If two mechanisms attract a similar distribution of bidders and also have similar “reserve prices,” the probability that at least one bidder has a value above the reserve price would be similar across the mechanisms.

Figure 4: Auctions and parcel leasing activity



For each quarter between 2004-2016, we compute the effect of auction status on the likelihood that a parcel has any active leases. We regress active status for parcel i during quarter t onto dummy equal to 1 for auction parcels and 0 for negotiable parcels, parcel covariates, and Grid 10 fixed effects.

$$\text{Leased}_{it} = \tau_t \text{Auction}_i + X_i \beta + \delta_{l(i)} + \epsilon_{it}$$

The figure shows point estimates and 95% confidence intervals for each quarter’s estimate of τ_t , with standard errors clustered at the grid level. A parcel is “Leased” if it has any active lease during a given quarter, including leases that were signed prior to 2004.

6.2 Parcel outcomes

We next look for differences in payment and production at the parcel level. As in Section 5, we restrict outcomes to those on leases signed between 2004 and 2013.³⁵ Table 7 presents the results. Like in our analysis of lease outcomes, we present two sets of models: linear regressions (per acre) and Poisson regressions. Across specifications, auction parcels earn about \$300 more per acre in discounted bonus payments in linear specifications and more than 60 log points more in Poisson specifications, and these differences are precisely estimated. The differences in output and lease revenues are all positive, and in most specifications are economically large. While these estimates are precisely estimated in some Poisson specifications, the estimates in linear specifications are quite noisy. This is likely due to the additional zeros that come from parcels with no leases, and the skewness of output and lease revenue, conditional on leasing. The combined effect of bonuses and royalty payments, reflected in the seller revenue models, suggests that auction parcels earn more than \$450 per acre than negotiation parcels. In the Poisson models, the difference is 30 log points, or more.

Table 7: Parcel Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Bonus	0.36 (0.08)	0.34 (0.10)	0.29 (0.04)	0.69 (0.08)	0.66 (0.10)	0.61 (0.06)
Auction - Output	0.09 (0.16)	0.21 (0.15)	0.16 (0.12)	0.22 (0.14)	0.29 (0.10)	0.16 (0.12)
Auction - Lease Revenue	0.23 (0.62)	0.58 (0.57)	0.63 (0.44)	0.20 (0.14)	0.26 (0.11)	0.13 (0.12)
Auction - Seller Revenue	0.46 (0.19)	0.52 (0.19)	0.45 (0.12)	0.37 (0.08)	0.40 (0.08)	0.30 (0.07)
Estimate	G10	G20	DML	G10	G20	DML
Estimator	Linear	Linear	Linear	Poisson	Poisson	Poisson
N	2,621	2,621	2,621	2,132	2,240	2,621

Columns 1-3 present ordinary least squares estimates of parcel level outcomes, *per acre*, while columns 4-6 show pseudo-Poisson maximum likelihood estimates of aggregate parcel outcomes in levels. Linear and Poisson fixed effect models include a spline in parcel size, while parcel size is part of the random forest covariates in both linear and Poisson DML models. Bonus measured in thousands of discounted dollars, BOE measured in hundreds of discounted barrels of oil equivalent, lease revenue measured in thousands of discounted dollars, and seller revenue in thousands of discounted dollars. All outcomes are discounted to January 1, 2004.

³⁵Although bonuses are not right censored, we restrict discounted bonuses here to leases signed by 2013 for consistency across parcel measures.

These parcel results confirm the main takeaways from Section 5: auction parcels pay lessors more and produce more output than negotiation parcels. However, the parcel results are generally smaller than the lease level results, and are considerably noisier. There are several reasons for this. First, many parcels never lease. While the previous section demonstrated that leasing was equally likely across parcel types, the parcel regressions still average in nearly thirty percent more zeros than the lease regressions. Second, the parcel regressions also include outcomes on leases that are excluded from the sample in Section 5. In our lease level analyses, we drop negotiated leases with features that do not overlap with auction leases, such as being too big or too small, or having multiple ownership interests. However, these lease level features are not defined at the parcel level, so we include these leases here. Third, even when leases are included in both samples, they may be weighted differently. When a lease spans two parcels, the outcomes from that transaction will get twice as much weight in a parcel level regression as a lease level regression. Conversely, if a single parcel has two leases on it, that parcel will effectively appear twice in the lease regressions but only once in a cross-sectional parcel regression. A final difference is the treatment of time. All lease regressions include quarter of sample fixed effects, to account for large changes in hydrocarbon prices and technology over time. However, the parcel regressions are inherently cross-sectional, with all outcomes discounted to 2004. In appendix A.5, we provide additional lease and parcel level regressions which remove some of these differences and provide more suitable comparisons. We find that a parcel-level regression restricted to parcels that ever lease yields results that are economically similar to, and often statistically indistinguishable from, a cross-sectional lease-level regression on the full population of leases.

7 Discussion

Auctioned leases pay and produce more than negotiated leases do. This is despite the fact that the land underlying auctioned and negotiated leases is nearly identical, as are the contracts governing the terms of investment and production. These differences in payment and output must therefore mean that auctions generate better matches than negotiations, in the sense that auctions allocate drilling rights to firms that value them more and can use them more productively. In this section, we discuss the potential sources of heterogeneity in match value for a given piece of land, and assess the plausibility that these difference could generate such large gains in payment and output.

Thousands of E&P companies operate in Texas, including more than 200 which produced more than 1,000 barrels of oil per day in 2020.³⁶ Within this group, there are wide differ-

³⁶For firm-level production statistics, see <https://www.rrc.state.tx.us/oil-and-gas/>

ences in size and observable measures of sophistication. Some of the largest publicly traded companies in the world, like Exxon and Chevron, compete for leases against a multitude of privately held E&P companies with fewer than 20 employees.³⁷ Beyond observable differences in firm size and sophistication, there is also evidence for heterogeneity across firms in their engineering designs of hydraulic fracturing treatments, which are necessary for all leases in this setting (Covert, 2015).

Given the many sources of vertical differentiation across firms, it is natural to ask whether auctions allocate to *different* firms. Table 8 provides an answer, by tabulating auction and negotiation market shares for the ten most active lessees in our sample.³⁸ It is noteworthy that all but one of these large firms are active in both mechanisms, winning at least one auction and several negotiations each.³⁹ Nevertheless, most firms have statistically and economically different shares of the auction and negotiation markets. The data soundly reject a χ^2 test of the hypothesis that a firm’s auction market share is the same as its negotiation market share ($p < 2 \times 10^{-16}$).⁴⁰ Thus, it does appear that the rate at which a firm wins a lease differs across auctions and negotiations.

To assess whether this differential firm composition across mechanisms explains our results, we compare auctions and negotiations *won by the same firm*. Table 9 shows estimates of our main bonus and output regressions, with and without fixed-effects for the identity of the firm that signed a lease. Bonus payments, output, and lease revenue are larger on auctioned leases than negotiated leases, regardless of whether we condition on the identity of the winning firm. In fact, point estimates for the within-firm gap between auction and negotiation outcomes are larger than than our baseline estimates for all three outcomes. These results show that the large differences in payment and production between auctions and negotiations are not just the result of differences in the composition of firms who win auctions relative to negotiations. Auctions must also generate better *horizontal* matches than negotiations do.

There are many potential sources of horizontal heterogeneity in valuations for a drilling

[research-and-statistics/operator-information/texas-oil-and-gas-producers-by-rank/](https://www.eia.doe.gov/data-graphics/operator-information/texas-oil-and-gas-producers-by-rank/)

³⁷For the firm size distribution in NAICS 20, Mining Quarrying & Oil and Gas Extraction, for Texas, see <https://www.census.gov/data/tables/2017/econ/susb/2017-susb-annual.html>

³⁸Firm identities in our data are recorded with some error (typos, etc). We describe our process for cleaning these names in Appendix B.

³⁹In further analysis of the auction bid data, we find that 62% of negotiated transactions are won by firms who place at least one bid at auction, and 75% of auctions are won by firms that successfully complete at least one negotiation.

⁴⁰ χ^2 tests of equal proportions for the top 20 and 40 most active lessees are similarly rejected. Moreover, these differences in market shares across the mechanism types do not simply reflect differences in the distribution of a firm’s “interest” across basins. We replicate this exercise within leases overlying the two largest shale basins in Texas, the Permian and the Eagle Ford, and can similarly reject a null hypothesis of equal proportions for the top 10 most active lessees in each basin.

Table 8: Auction and Negotiation Market Shares Among Active Lessees

Firm	Leases	Auction Share	Negotiation Share	Overall Share
CHESAPEAKE	125	0.161	0.049	0.083
PETRO HUNT	98	0.007	0.090	0.065
ENERGEN	92	0.062	0.060	0.061
PETROHAWK	72	0.084	0.032	0.048
SEGUNDO NAVARRO	65	0.004	0.059	0.043
ANADARKO	61	0.042	0.040	0.040
BP	59	0.000	0.056	0.039
CIMAREX	58	0.040	0.038	0.038
777	36	0.011	0.029	0.024
DEVON	33	0.059	0.006	0.022
ALL OTHERS	816	0.531	0.542	0.539

Table 9: Effects of Firm Composition and Mechanism Type on Lease Outcomes

	Bonus	Bonus	Output	Output	Revenue	Revenue
Auction	0.53 (0.06)	0.58 (0.08)	0.821 (0.403)	1.583 (0.681)	2.64 (1.46)	4.67 (2.07)
Firm FE	No	Yes	No	Yes	No	Yes
N	1,515	1,515	1,259	1,259	1,259	1,259
R^2	0.961	0.974	0.742	0.828	0.765	0.857

The dependent variable is the natural logarithm of the bonus payment (columns 1 and 2), discounted barrels of oil equivalent per acre (columns 3 and 4), or discounted lease revenue per acre (columns 5 and 6). In columns 2-6, the sample includes all leases whose primary term ends before March, 2019. All specifications include fixed effects for 10-mile grids-by-year and quarter-of-sample, as well as a spline in acres.

opportunity across firms, in the sense that firm A might be able to use lease X more productively than firm B, while B might be able to use lease Y more productively than firm A. For example, firms who already control acreage in the area may be able to develop a more efficient drilling plan. Firms who own hydrocarbon transportation infrastructure close to a given parcel may experience cost advantages in developing that specific parcel. Similarly, firms with formation-specific knowledge about geology or efficient engineering choices will be able to produce more (or less expensively) than firms with less context-specific knowledge.

We can use auction bid data to measure the importance of such horizontal match quality across firm-lease pairs. GLO reports the bids submitted by all auction participants, not just the winner. Among large lessees, it is possible to observe the same two firms bidding against each other in many auctions. If lease valuations were mostly vertical, in the sense that one firm was inherently more productive on all parcels than another, then the more productive firm would reliably bid more than the less productive firm. In contrast, if lease valuations were mostly driven by horizontal match quality, bid rankings would be less predictable. We check this in the bid data, by looking at all pairs of firms who bid in the same auction ten or more times. Table 10 lists these pairs and tabulates the probability that the alphabetically earlier firm (Firm A) bids more than the later firm (Firm B). For 10 of the 12 pairs of firms, the fraction of the time that one firm wins more than the other is statistically identical to a coin toss, which strongly rejects the hypothesis that values are mostly vertical, at least among this set of firms. To the extent that there are persistent vertical differences between firms, this data suggests that those are small relative to horizontal differences in firm-lease match values.

The previous discussion lays out several sources of heterogeneity in preferences for a lease across firms, and provides evidence that they exist in our sample. Are these differences large enough to rationalize the large treatment effects we estimate? We can again use the auction bid data to answer this, using techniques from modern empirical auction analysis.

Table 10: Bid ranking for top auction pairs

Firm A	Firm B	Auctions	Share A > B	p-value
CIMAREX	ENERGEN	31	0.52	1.000
CIMAREX	CONOCOPHILIPS	19	0.79	0.019
CIMAREX	RESOLUTE	19	0.53	1.000
CONOCOPHILIPS	ENERGEN	19	0.37	0.359
ENERGEN	RESOLUTE	19	0.42	0.648
COG	RANGE	17	0.41	0.629
CONOCOPHILIPS	RESOLUTE	17	0.53	1.000
CIMAREX	MARSHFIELD OIL AND GAS LLC	12	0.67	0.388
ENERGEN	MARSHFIELD OIL AND GAS LLC	12	0.67	0.388
ADESCAPE INC	ENERGEN	10	0.00	0.002
ADESCAPE INC	SEMEION RESOURCES LP	10	0.30	0.344
ENERGEN	SEMEION RESOURCES LP	10	0.80	0.109

p -value from a two-sided exact test of the hypothesis that the share of auctions in which Firm A bids more than Firm B is equal to 0.5, at a 95% level.

Table 11: Auction Results by Number of Bidders

Bids	Auctions	Average Reserve	Average Bonus	Average Reserve Margin	Average 1-to-2 Markup	Average 1-to-N Markup	Average 1-to-2 Allocative Gain	Average 1-to-N Allocative Gain
1	252	704	1,325	0.54				
2	97	860	1,790	0.71	0.38		0.50	
3+	105	882	4,191	1.34	0.30	0.86	0.43	1.07

This table summarizes bids from GLO auctions. The Bonus is the winning bid, in dollars per acre. The “Reserve Margin” is the natural logarithm of the ratio of the winning bid to the reserve price. The “1-to-2 Markup” is the natural logarithm of the ratio of the winning bid to the second highest bid, and the “1-to-2-N Markup” is the same, but comparing the highest and lowest bids. The “1-to-2 Allocative Gain” is the natural logarithm of the ratio of the GPV pseudovalues of the winning and second highest bids, and “1-to-N Allocative Gain” is the same, but comparing the GPV pseudovalues of the highest and lowest bids.

Table 11 summarizes the bid data from all 454 GLO leases in our sample which cleared above the reserve. More than half of auctions receive only a single bid above the posted reserve price, and fewer than a quarter receive three or more bids.⁴¹ A key fact in this data is that winning bids are *substantially* higher than non-winning bids, and regardless of how

⁴¹During this time period, GLO auctions receive somewhat fewer bids per lease than neighboring New Mexico’s state mineral auction does. The average successful auction in this sample has 2 bids. In comparison, [Bhattacharya et al. \(2022\)](#) and [Kong \(2020\)](#) report 2.4 - 2.6 bids per auction in New Mexico.

many bids an auction receives, those bids are quite a bit higher than reserve, even when there is no competition for the winner, *ex post*.⁴² Moreover, winning bids are more than a little bit higher than losing bids. For auctions with two or more bidders, the average winning bid is about 34 log points higher than the second highest bid, with smaller differences in auctions with three or more bidders. The gap between winning bids and lower bids is even larger. For auctions with three or more bidders, the difference between the winning bid and the *lowest* bid is 86 log points, on average. Thus, even when more than two bidders arrive at an auction, there are enormous differences in their apparent willingness-to-pay for the same property right.

We can compare moments from the distribution of auction bids to our empirical estimates from sections 5 and 6. The average difference between negotiated and auction bonus payments in Table 4 is about 50 log points. This is close to the average reserve margin for auctions with just one bidder (54 log points), but smaller than the average across all auctions (76 log points). The highest 95% upper confidence interval for the difference between auction and negotiation lease bonuses in Table 4 is only about 68 log points, so the average gap between auctions and negotiations is clearly smaller than the reserve margin. However this gap is larger than the average difference in bids between the highest and second highest bidder.⁴³

We can make this comparison between negotiated bonuses and losing auction bids explicit using regression. Table 12 repeats the log bonus regressions from Section 5, but replaces an auction’s winning bid, which is its true bonus payment, with other bids placed for that lease (we do not modify negotiated bonus payments). The first two columns recreate the main results from Table 4, showing that auctioned leases earn bonus payments that are 53-54 log points larger than similar negotiated bonus payments. In columns 3 and 4, we replace realized auction bonuses with their reserve prices. On average, auction reserve prices are somewhat *lower* than negotiated bonus payments on leases that transact at similar places and times, which means that the core bonus payment result cannot be explained by negotiated leases frequently transacting at prices below the minimum level allowed in nearby auctions. Columns 5 and 6 take this one step further, replacing auction bonuses with the

⁴²Kong (2020) also observes this pattern in mineral lease auctions which take place in neighboring New Mexico, where the average single bidder auction transacts at *seven* times the reserve price. Kong demonstrates that this behavior can be rationalized by bidder’s having uncertainty about how much competition will show up to an auction.

⁴³Note that these differences are constructed from *within lease* comparisons, winning vs. losing bids for the same contract on the same piece of land, at the same point in time. This difference is somewhat closer to the estimates in Table A.4, which compares auctions and negotiations that are both similar in location and time, *and* have similar royalty rates and primary terms.

minimum bid received.⁴⁴ In this comparison, we see that negotiations earn bonus payments that are similar, on average, to the lowest payment in nearby auctions. Finally, columns 7 and 8 replace the auction bonus with the second highest bid. Overall, these results are consistent with the results in Table 11. Negotiated transactions occur at prices above the reserve prices of auctions for similar leases, but at or slightly below the second highest bid in those auctions.

Table 12: $\log(\text{Bonus})$ results under alternative auction payment assumptions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Auction	0.53 (0.06)	0.58 (0.05)	-0.14 (0.04)	-0.07 (0.04)	0.00 (0.05)	0.06 (0.04)	0.10 (0.06)	0.15 (0.05)
Auction Bonus	Winner	Winner	Reserve	Reserve	Worst	Worst	Second	Second
Grid	10	DML	10	DML	10	DML	10	DML
Time	GY,Q	DML	GY,Q	DML	GY,Q	DML	GY,Q	DML
N	1,515	1,515	1,515	1,515	1,515	1,515	1,515	1,515

“Winner” means auction leases get their actual bonus payment (the winning bid). “Reserve” means auction leases get the reserve price instead of the winning bid. “Worst” means auction leases get the worst price among all submitted bids, or the reserve price if only 1 bid was submitted. “Second” means auction leases get the second best price among all submitted bids, or the reserve price if only 1 bid was submitted. Fixed effect specifications include a spline in acres, while DML specifications include acres in the random forest component.

We can infer the underlying lessee values associated with these auction bids using a procedure similar to Roberts (2013).⁴⁵ In the second to last column of Table 11, we compare the winner’s value to the second highest bidder’s value. As in our analysis of the bids themselves, we see that there are large differences between winning and losing values. For auctions with two bidders, the winner’s value is 50 log points higher than the second highest bidder’s value is, on average, and while the gap is smaller for auctions with more bidders, which are about half the sample, it is still economically large in auctions with three or more bidders, at 43 log points. The difference between highest and lowest values, in auctions with three or more bidders, is even larger, at 107 log points.

It is somewhat harder to compare the distribution of bidder values with our output results. Bidders care about production that is expected to occur long after our royalty revenue data ends, so realized production may provide a negatively biased measure of bidder value. On the other hand, bidders, as potential operators of mineral leases, also care about the costs of developing those leases, which we cannot measure with our output data. Having

⁴⁴For auctions with 1 bid, we define the minimum bid as the reserve price.

⁴⁵See Appendix D for implementation details.

noted those caveats, we compare the estimated allocative gains in the auction data with the results in Table 6. The 1-2 allocative gain, which is the average difference in log bidder values between the winner and second highest bidder, is well within the range of the point estimates in bottom panel of Table 6. Thus, an auction which mistakenly allocated a lease to the second highest bidder would suffer a decrease in allocative efficiency that is about as large as the differences in output that we estimate between similar auctioned and negotiated leases. However, the estimated output effects are considerably smaller than the average 1-N allocative gain.

7.1 Mechanism

Mineral leases that are allocated by auction are better matched to E&P firms than those allocated via negotiation. What features of negotiation process lead to worse matches? Unfortunately, we cannot answer this question directly because there are no records of the circumstances that lead up to a successful negotiated transaction, nor are there any records of initiated but failed negotiations. In lieu of sufficient transaction level detail to quantitatively evaluate the negotiation process, we instead discuss how institutional features of this market and the resulting differences in outcomes fit within existing mechanism comparisons considered by the literature.

The calculations in the previous section tell us how much lower negotiation outcomes would be if negotiators effectively faced the same set of bidders as auctions, but mistakenly (and unexpectedly to bidders) awarded the lease to the wrong party. A more tenable mechanism would be that RAL negotiators actually use a similarly formal and simultaneous mechanism, but these RAL “auctions” simply attract fewer bidders than GLO auctions do. This is roughly the “non-sequential” search mechanism considered by [Salz \(2022\)](#). In our setting, increased demand for GLO auctions, relative to hypothetical RAL auctions would arise from the fact that they are centralized. State auctions are widely publicized, routinely held, and operate under known reserve prices, whereas a central challenge for firms in acquiring negotiated acreage (both in RAL and private land writ large) is identifying which land is leasable, performing title research to determine who actually owns it, and estimating the surface owner’s latent willingness-to-sell. It is thus likely that the latter mechanisms would result in fewer participants.

We can use our bidder value estimates to evaluate the allocative efficiency costs of simply having access to fewer bidders in a standard auction framework. To do this, we simulate $v_{(i:k)}$, the expected value of the natural logarithm of the i -th highest value out of k draws from the distribution of estimated values and compare this value for different values of i and k . With

2 bidders, losing a bidder lowers the average value of the winner by $v_{(1:2)} - v_{(1:1)} \approx 0.24$. For three bidders, the losses are smaller, approximately 0.15. Since few auctions attract more than 3 bidders, it is unlikely that the gap in output we observe between auctions and negotiations is *only* reflective of a gap in competition between the two mechanisms, as these differences in value are smaller than the differences observed in Table 6. It is possible to construct differences closer to those in Table 6, but they require dropping more than 1 bidder. For example, going from 3 bidders to just 1 generates a difference of $v_{(1:3)} - v_{(1:1)} \approx 0.39$.

What little we know about the private leasing process suggests that the mechanism is more sequential than simultaneous. The theory literature offers conflicting opinions as to whether a sequential mechanism will perform better than a simultaneous auction. If participation is costly, a sequential mechanism saves real resources by allowing bidders to observe existing bids before deciding to incur entry costs. Bulow and Klemperer (2009) show that while a sequential mechanism is thus always more efficient, an auction is *usually* more profitable for the seller. However, the latter result is predicated on their assumption that a bidder’s entry choice is independent of its value for the lease. Roberts and Sweeting (2013) demonstrate that a similar sequential mechanism can outperform auctions if this entry choice is instead *selective*, in the sense that better users of a lease are more likely to participate than worse users. Thus, if the *only* difference between the informal process for RAL negotiations and the GLO’s auctions was that auctions considered bids simultaneously, while negotiations reviewed offers from the same set of bidders sequentially, then the increased revenues and output that auctions generate in our setting suggests that negotiation entry choices by E&P companies are not especially selected.

While it is possible to rationalize our empirical results by assuming that negotiations are simply auctions with fewer bidders or by assuming that negotiations follow a formal sequential process, it is important to note that neither assumption perfectly fits this setting. In the primary market for oil and gas leases, offers to mineral owners are initiated by the buyer. Savvy leasing agents, cognizant of the relative unsophistication of their counterparts, likely use a variety of persuasive techniques which do not fit well within a formal mechanism design framework. Conversely, what little information is available suggests that private landowners are not particularly savvy. In the most extensive survey of private mineral rights owners to date, 79% of lessors in the highly active Marcellus Shale region report speaking to only one E&P company before signing a lease.⁴⁶ They also appear relatively uninformed, with only 32% reporting to have consulted any educational materials prior to signing.⁴⁷

⁴⁶Survey conducted by the Penn State Extension Marcellus Education Team and summarized in “Natural Gas Lessors’ Experiences in Bradford and Tioga Counties, 2010” [Online version available [here](#), accessed 3/15/2018].

⁴⁷In Appendix Table A.8, we explore the hypothesis that most lessors are “unsophisticated” by including

Relatedly, it seems intuitive that landowners would have a difficult time committing to (and executing) a more formal process. In the Pennsylvania survey, only 21% of lessors reported ever consulting with a lawyer before transacting.

Another possible explanation for the large differences we observe between auctions and negotiations is that lessees are colluding more in the latter market. Competition authorities have previously investigated bid-rigging in both mineral lease auctions⁴⁸ and mineral lease negotiations.⁴⁹ Theoretically, it is unclear whether auctions and negotiations should be differentially susceptible to collusion, and we are not aware of any empirical evidence on this question. Collusive behavior usually requires bidders who can monitor each other's behavior, and this may be easier to do in auctions than negotiations. GLO publishes the identity of all bidders and the bids they make in each auction, while there is no record of what happens during the negotiation process, aside from who the ultimate lessee is and what price they paid. However, the detection of collusive behavior by regulators may also be easier in auctions, as a result of this publicly disclosed data.

7.2 External Validity

We conclude by considering how generalizable the results in this paper are to the broader population of mineral leases on private land in the United States, which are also allocated in an informal, decentralized fashion. On the one hand, Relinquishment Act lessors are likely better informed than the general private mineral rights owner population. Although the process for RAL leasing initially mirrors that of private leasing, with a landman approaching the surface owner with an offer and the two parties coming to a private agreement, these agreements must be approved by the GLO before they are finalized. During this approval process, the terms of the agreement may be improved, with the GLO requesting, for example, a higher bonus payment or royalty rate. In our sample, 19% of RAL leases show some type of improvement during this approval period: the median improvements for bonuses and royalties are 50% and 17%, respectively. Throughout this paper, we compare realized lease terms from RAL negotiations, rather than what the landowners would have negotiated absent state intervention. If RAL surface owners had only acted on their own, it is likely that we'd find even larger differences between auctioned and negotiated outcomes than we currently

measures of lessor experience and sophistication in our main bonus regressions. In this analysis, we find suggestive evidence that more sophisticated RAL lessors, those that sign many leases, are corporations, or are directly involved in the oil & gas business, secure moderately higher bonus payments than other RAL lessors do. However, these differences are still much smaller than the difference between auctions and negotiations.

⁴⁸For an example involving *federal* mineral lease auctions, see <https://www.wilmerhale.com/en/insights/publications/us-justice-department-announces-first-ever-settlement-of-false-claims-act-and-an>

⁴⁹See <https://www.reuters.com/article/us-chesapeake-encana-antitrust-idUSBRE91015D20130225>.

do.

On the other hand, a major difference between the negotiations we observe in this setting and fully private mineral lease negotiations is that RAL surface owners only enjoy half of the economic benefits of leasing, but experience all of the costs of getting a lease negotiated in the first place. This agency problem suggests that RAL surface owners might have insufficient incentives to negotiate good leases, and that lessors with a full stake in the returns to their efforts might perform closer to the auction outcome. While this is conceptually plausible, we doubt that it explains much of our results, for several reasons. First, in the aforementioned Pennsylvania survey, most lessors report very little search effort or success, despite having no agency issues. Second, in this setting the principal (GLO) observes agent effort, and can effectively require surface owners to search more or negotiate harder if it believes it is possible. Third, we test for evidence of shirking by comparing RAL performance on large vs small leases, and find that they do similarly poorly relative to auctions. As we discuss in Appendix A.4, because negotiating effort is largely a fixed cost, if shirking were an important part of the story, then we would expect RAL surface owners with larger parcels to negotiate higher bonuses than those with smaller parcels. For these reasons, we believe the agency problems that are unique to this setting do not limit the validity of these results to other settings. Negotiators perform poorly not because they are shirking, but because they are uninformed and/or unskilled.

Finally, we want to emphasize that the auction “counterfactual” we have in mind is easy to imagine. Although we study a setting where auctions are administered by a state agency, auction technology is generally available to all private mineral owners. While the costs associated with organizing an auction may have been large prior to the Internet era, today there are electronic mineral auction platforms whose fees are 10% or less of the final transaction price. Indeed, the GLO now uses one such platform, EnergyNet.com, that explicitly advertises its availability to private landowners. Given how large the difference between auctioned and negotiated payments are, the gain from using an auction appears to far exceed the cost.⁵⁰ In this specific context, it’s also possible to imagine GLO performing these auctions on the surface owner’s behalf, and presumably internalizing some scale economies while doing so.⁵¹

⁵⁰Since RAL landowners only have a 50% claim to the gain from auctions, the effective fee from the RAL landowners perspective would be 20%, which is still far below the estimated auction gain.

⁵¹Indeed, GLO already does this when E&P firms wish to lease minerals in RAL parcels in which ownership cannot be established, due to inheritance or property title issues.

8 Conclusion

At current prices, proved US oil and gas reserves are worth approximately \$4.5 trillion, and the vast majority of these resources are owned, and managed, by private individuals. While this arrangement has delivered substantial wealth to countless landowners, the informal mechanisms they use to find and bargain with their contracting partners may generate less revenue and less efficient matches to E&P companies than would be possible under a more formal mechanism. In this paper, we directly quantify this loss. Using rich data on a large number of leases affected by a natural experiment, we compare outcomes under unstructured “negotiations” to formal auctions. Our results show that auctions generate 53 more log points in up front payments, and that auctions produce 39 more log points in output, suggesting that auctions facilitate better matches between land and the firms that can use it most productively. Given that landowners in this setting often have assistance from an informed third party in the form of the Texas General Land Office, these results likely provide a lower bound on the prospective gains from using auctions in the private mineral leasing population writ large.

A natural direction for future work would be to investigate why informal mechanisms perform so poorly. In this paper, we lack sufficient information on the process leading up to informal transactions, and instead rely on credible identification of the net effect of formal vs. informal mechanisms in the “reduced form.” One approach to gaining insight about the causes of this difference would be to perform surveys of informal mechanism users or to conduct experimental information interventions, in mineral leasing or other settings. Another would be to measure similar reduced form differences in other economically important markets where formal and informal mechanisms coexist, such as real estate, construction procurement, and used automobile sales. In these other settings, sellers may be more or less informed, or have different abilities to attract potential buyers. Given the sheer size of these other markets, if even a fraction of the estimated gains in this paper translate, the gains from policy that encourages the use of formal mechanisms would be enormous.

References

- Allen, Jason, Robert Clark, and Jean-François Houde**, “Search frictions and market power in negotiated-price markets,” *Journal of Political Economy*, 2019, 127 (4), 1550–1598.
- Angrist, Joshua D and Jörn-Steffen Pischke**, *Mostly harmless econometrics: An empiricist’s companion*, Princeton University Press, 2008.

- Athey, Susan, Julie Tibshirani, Stefan Wager et al.**, “Generalized random forests,” *The Annals of Statistics*, 2019, *47* (2), 1148–1178.
- Backus, Matthew, Thomas Blake, and Steven Tadelis**, “On the empirical content of cheap-talk signaling: An application to bargaining,” *Journal of Political Economy*, 2019, *127* (4), 1599–1628.
- , – , **Brad Larsen, and Steven Tadelis**, “Sequential bargaining in the field: Evidence from millions of online bargaining interactions,” *The Quarterly Journal of Economics*, 2020, *135* (3), 1319–1361.
- Bajari, Patrick, Robert McMillan, and Steven Tadelis**, “Auctions Versus Negotiations in Procurement: An Empirical Analysis,” *The Journal of Law, Economics, and Organization*, October 2009, *25* (2), 372–399.
- Bartik, Alexander W, Janet Currie, Michael Greenstone, and Christopher R Knittel**, “The local economic and welfare consequences of hydraulic fracturing,” *American Economic Journal: Applied Economics*, 2019, *11* (4), 105–55.
- Bhattacharya, Vivek, Andrey Ordin, and James W Roberts**, “Bidding and drilling under uncertainty: An empirical analysis of contingent payment auctions,” *Journal of Political Economy*, 2022, *130* (5), 1319–1363.
- Boone, Audra L. and J. Harold Mulherin**, “How Are Firms Sold?,” *The Journal of Finance*, March 2007, *62* (2), 847–875.
- Brehm, Paul A and Eric Lewis**, “Information asymmetry, trade, and drilling: evidence from an oil lease lottery,” *The RAND Journal of Economics*, 2021, *52* (3), 496–514.
- Brown, Jason P., Timothy Fitzgerald, and Jeremy G. Weber**, “Capturing rents from natural resource abundance: Private royalties from U.S. onshore oil & gas production,” *Resource and Energy Economics*, November 2016, *46*, 23–38.
- Bulow, Jeremy and Paul Klemperer**, “Auctions Versus Negotiations,” *The American Economic Review*, 1996, *86* (1), 180–194.
- and – , “Why Do Sellers (Usually) Prefer Auctions?,” *The American Economic Review*, 2009, *99* (4), 1544–1575.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins**, “Double/debiased machine learning for treatment and structural parameters,” *The Econometrics Journal*, 01 2018, *21* (1), C1–C68.
- , **Ivan Fernandez-Val, and Alfred Galichon**, “Improving point and interval estimators of monotone functions by rearrangement,” *Biometrika*, 2009, *96* (3), 559–575.
- Covert, Thomas**, “Experiential and Social Learning in Firms: the Case of Hydraulic Fracturing in the Bakken Shale,” 2015. working paper.

- Cuesta, José Ignacio and Alberto Sepulveda**, “Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access,” *Available at SSRN 3282910*, 2018.
- Currie, Janet, Michael Greenstone, and Katherine Meckel**, “Hydraulic fracturing and infant health: New evidence from Pennsylvania,” *Science Advances*, December 2017, *3* (12), e1603021.
- Farrell, Joseph**, “Information and the Coase Theorem,” *The Journal of Economic Perspectives*, December 1987, *1* (2), 113–129.
- Feyrer, James, Erin T. Mansur, and Bruce Sacerdote**, “Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution,” *American Economic Review*, April 2017, *107* (4), 1313–1334.
- Fitzgerald, Timothy and Randal Rucker**, “US private oil and natural gas royalties: estimates and policy relevance,” *OPEC Energy Review*, March 2016, *40* (1), 3–25.
- Gentry, Matthew and Caleb Stroup**, “Entry and competition in takeover auctions,” *Journal of Financial Economics*, 2018.
- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong**, “Optimal nonparametric estimation of first-price auctions,” *Econometrica*, 2000, *68* (3), 525–574.
- Heckman, James J**, “Sample selection bias as a specification error,” *Econometrica: Journal of the econometric society*, 1979, pp. 153–161.
- Hendricks, Kenneth and Robert H. Porter**, “An Empirical Study of an Auction with Asymmetric Information,” *The American Economic Review*, 1988, *78* (5), 865–883.
- and –, “The Timing and Incidence of Exploratory Drilling on Offshore Wildcat Tracts,” *The American Economic Review*, June 1996, *86* (3), 388–407.
- Herrnstadt, Evan, Ryan Kellogg, and Eric Lewis**, “Royalties and Deadlines in Oil and Gas Leasing: Theory and Evidence,” 2022. working paper.
- Hortaçsu, Ali and Chad Syverson**, “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds,” *The Quarterly Journal of Economics*, 05 2004, *119* (2), 403–456.
- Kong, Yunmi**, “Not knowing the competition: evidence and implications for auction design,” *The RAND Journal of Economics*, 2020, *51* (3), 840–867.
- , “Sequential auctions with synergy and affiliation across auctions,” *Journal of Political Economy*, 2021, *129* (1), 000–000.
- Larsen, Bradley J**, “The efficiency of real-world bargaining: Evidence from wholesale used-auto auctions,” *The Review of Economic Studies*, 2021, *88* (2), 851–882.

- Leonard, Bryan and Dominic P Parker**, “Fragmented Ownership and Natural Resource Use: Evidence from the Bakken,” *The Economic Journal*, 2021, 131 (635), 1215–1249.
- Muehlenbachs, Lucija, Elisheba Spiller, and Christopher Timmins**, “The Housing Market Impacts of Shale Gas Development,” *American Economic Review*, December 2015, 105 (12), 3633–3659.
- Myerson, Roger B and Mark A Satterthwaite**, “Efficient mechanisms for bilateral trading,” *Journal of economic theory*, 1983, 29 (2), 265–281.
- P2 Energy Solutions**, “Texas Permanent School Fund Land Grid,” February 2018.
- Roberts, James W**, “Unobserved heterogeneity and reserve prices in auctions,” *The RAND Journal of Economics*, 2013, 44 (4), 712–732.
- Roberts, James W. and Andrew Sweeting**, “When Should Sellers Use Auctions?,” *American Economic Review*, August 2013, 103 (5), 1830–1861.
- Robinson, Peter M**, “Root-N-consistent semiparametric regression,” *Econometrica: Journal of the Econometric Society*, 1988, pp. 931–954.
- Salz, Tobias**, “Intermediation and competition in search markets: An empirical case study,” *Journal of Political Economy*, 2022, 130 (2), 310–345.
- Shields, James D.**, “Leasing Lands Subject to the Texas Relinquishment Act Comment,” *St. Mary’s Law Journal*, 1981, 13, 868–886.
- Subramanian, Guhan**, “Go-Shops vs. No-Shops in Private Equity Deals: Evidence and Implications,” *Business Lawyer (ABA)*, 2007, 63, 729.
- Syverson, Chad**, “What determines productivity?,” *Journal of Economic literature*, 2011, 49 (2), 326–65.
- Texas Department of Transportation**, “TxDOT Roadways,” August 2017.
- Texas General Land Office**, “Past Bid Sale Results,” April 2017.
- , “Active Oil & Gas Leases,” June 2020.
- , “Inactive Oil & Gas Leases,” June 2020.
- Timmins, Christopher and Ashley Vissing**, “Environmental justice and Coasian bargaining: The role of race, ethnicity, and income in lease negotiations for shale gas,” *Journal of Environmental Economics and Management*, 2022, p. 102657.
- U.S. Energy Information Administration**, “Eagle Ford play boundaries, structure and isopachs,” May 2019.
- , “Henry Hub Natural Gas Spot Price,” September 2019.
- , “Low permeability oil and gas play boundaries in Lower 48 States,” May 2019.

– , “Permian Basin: Wolfcamp formation elevation and isopachs,” May 2019.

– , “West Texas Intermediate Crude Oil Spot Price,” September 2019.

U.S. Geological Survey, “National Hydrography Dataset,” October 2019.

– , “National Land Cover Database,” June 2021.

Vissing, Ashley, “One-to-Many Matching with Complementary Preferences: An Empirical Study of Market Power in Natural Gas Leasing,” August 2017.

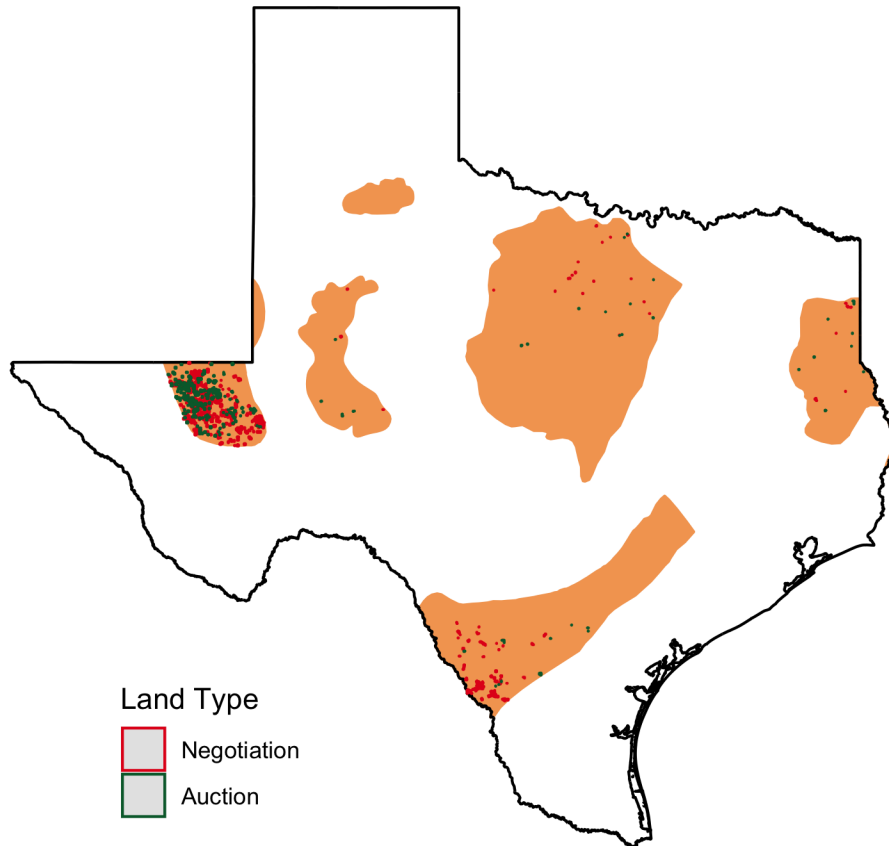
Xie, Jun and Chaofeng Liu, “Adjusted Kaplan–Meier estimator and log-rank test with inverse probability of treatment weighting for survival data,” *Statistics in medicine*, 2005, *24* (20), 3089–3110.

Yergin, Daniel, *The Prize* December 2008.

Appendix A Additional Tables and Figures

A.1 RAL vs State Lease Locations

Figure A.1: Map of Sample Leases by Type



A.2 Additional Lease Results

Table A.1: Bonus Payments and Mechanism Type, per Acre

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Auction	0.63 (0.14)	0.66 (0.18)	0.82 (0.28)	0.64 (0.19)	0.78 (0.11)	0.66 (0.19)	0.77 (0.10)	0.87 (0.34)	1.06 (0.19)
Grid	10	10	10	20	DML	10	DML	10	DML
Time	Q	GY,Q	GYQ	GY,Q	DML	GY,Q	DML	GY,Q	DML
Extra	No	No	No	No	No	Yes	Yes	No	No
Private Only	No	No	No	No	No	No	No	Yes	Yes
N	1,515	1,515	1,515	1,515	1,515	1,260	1,260	1,308	1,308
R^2	0.728	0.880	0.901	0.807		0.868		0.890	

The dependent variable in each regression is the the lease’s bonus payment per acre. In columns 1-4, the size of the location bins, in miles, are indicated in the “Grid” row, while the structure of the time controls (“Q” for quarter of sample, “GY,Q” for grid-by-year plus quarter of sample, and “GYQ” for grid-by-quarter of sample) are indicated in the “Time” row. Standard errors are clustered by grid in columns 1-4. Column 5 uses a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). Fixed effect models include a spline in lease size while the DML model includes lease size as a random forest covariates.

Table A.2: Lease Bonus Regressions with Controls for Royalty Rate and Term

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.72 (0.18)	0.68 (0.16)	0.96 (0.13)	0.43 (0.06)	0.43 (0.05)	0.57 (0.05)
Estimate	G10Y	G20Y	DML	G10Y	G20Y	DML
Outcome	Bonus per acre	Bonus per acre	Bonus per acre	log(Bonus)	log(Bonus)	log(Bonus)
N	1,515	1,515	1,515	1,515	1,515	1,515

The dependent variable is either bonus per acre (in thousands per acre) or the natural logarithm of the bonus payment. The Grid10-Year and Grid20-Year fixed-effect specifications include a spline in lease acreage and linear terms for term length and royalty rate, while the DML specifications include these variables as random forest covariates.

Table A.3: Lease Output Regressions with Controls for Royalty Rate and Term

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Lease Revenue	3.85 (1.46)	4.44 (1.54)	6.14 (1.90)	0.53 (0.20)	0.77 (0.22)	0.75 (0.08)
N	1,259	1,259	1,259	746	980	1,259
Auction - Output	1.17 (0.41)	1.32 (0.41)	1.55 (0.43)	0.62 (0.21)	0.83 (0.25)	0.81 (0.24)
N	1,259	1,259	1,259	746	980	1,259
Auction - Seller Revenue	1.31 (0.39)	1.56 (0.39)	2.16 (0.49)	0.48 (0.12)	0.65 (0.14)	0.84 (0.23)
N	1,259	1,259	1,259	1,259	1,259	1,259
Estimate	G10Y	G20Y	DML	G10Y	G20Y	DML
Estimator	Linear	Linear	Linear	Poisson	Poisson	Poisson

The dependent variables are per-acre (lease and seller revenue in thousands per acre, output in hundreds of BOE per acre) in the linear specifications and in aggregate (not per acre) in the Poisson specifications. The Grid10-Year and Grid20-Year fixed-effect specifications include a spline in lease acreage and linear terms for term length and royalty rate, while the DML specifications include these variables as random forest covariates.

Table A.4: Drilled Regressions

	(1)	(2)	(3)	(4)	(5)
Auction	0.07 (0.03)	0.05 (0.04)	0.08 (0.06)	0.07 (0.03)	0.04 (0.03)
Grid	10	10	10	20	DML
Time	Q	GY,Q	GYQ	GY,Q	DML
N	1,259	1,259	1,259	1,259	1,259
R^2	0.375	0.626	0.719	0.489	

The dependent variable is equal to 1 if a lease was drilled during its primary term and 0 otherwise. The sample includes all leases whose primary term ends before March, 2019. In columns 1-4, the size of the location bins, in miles, are indicated in the “Grid” row, while the structure of the time controls (“Q” for quarter of sample, “GY,Q” for grid-by-year plus quarter of sample, and “GYQ” for grid-by-quarter of sample) are indicated in the “Time” row. Standard errors are clustered by grid in columns 1-4. Column 5 uses a double/debiased machine learning routine. Fixed effect models include a spline in lease size and the DML model includes lease size as a random forest covariate.

Table A.5: Log Output Regressions Among the Sample of Leases That Are Drilled

	(1)	(2)	(3)	(4)	(5)
Auction	0.40 (0.19)	0.52 (0.20)	0.33 (0.25)	0.58 (0.22)	0.45 (0.17)
Grid	10	10	10	20	DML
Time	Q	GY,Q	GYQ	GY,Q	DML
N	425	425	425	425	425
R^2	0.652	0.856	0.909	0.774	

The dependent variable is the natural logarithm of discounted barrels of oil equivalent (Output) per acre. The sample includes all drilled leases whose primary term ends before March, 2019. In columns 1-4, the size of the location bins, in miles, are indicated in the “Grid” row, while the structure of the time controls (“Q” for quarter of sample, “GY,Q” for grid-by-year plus quarter of sample, and “GYQ” for grid-by-quarter of sample) are indicated in the “Time” row. Standard errors are clustered by grid in columns 1-4. Column 5 uses a double/debiased machine learning routine. Fixed effect models include a spline in lease size and the DML model includes lease size as a random forest covariate.

A.3 Unleased Spell Duration Analysis

Our lease and parcel data make it possible to reliably identify which parcels have active leases starting in January, 2001. We leases to their parcels using our existing parcel-lease map and then merge this with our lease revenue data, which starts in January, 2001. We define a lease as active during a given month if that month is during the lease’s primary term or if that month is spanned by the earliest and latest royalty payments we observe for the lease. We say a lease is terminated at the later of a lease’s primary term expiration and its last observed royalty payment. If a given terminated lease is the only lease currently active in a parcel, we say that parcel is *unleased* at that point. As soon as another lease is signed that is associated with a parcel, we say the parcel as *leased* again. With this structure, every parcel in our data is either leased or unleased on January 1, 2001, and for every subsequent month, and we use this parcel-by-month panel to construct “unleased spells” for every parcel. It is worth noting that parcels which are unleased as of January 1, 2001 may have been unleased previous to this point. As a result, unleased spells that begin in January, 2001 may be left-censored.

In the unleased spell analyses that follow, we define four samples of unleased spells. The “All” spell sample is everything described above, while the “First” spell sample is just the first unleased spell for each parcel. The “Uncensored” spell sample excludes unleased spells that begin in January, 2001, because we do not know whether they began earlier than January 2001 or not. Finally, the “First Uncensored” spell sample is the first uncensored unleased spell for each parcel. Table A.6 tabulates how many spells we have under each of these definitions, for RAL and State parcels, starting in each year between 2001 and 2019.

Our first test of the hypothesis that RAL and State parcels experience differentially long unleased spells makes use of the adjusted Kaplan-Meier survival curve estimator, derived in


Table A.6: Parcel Unleased Spell Counts

Year	All		Uncensored		First		First Uncensored	
	RAL	State	RAL	State	RAL	State	RAL	State
2001	1439	450	17	10	1439	450	17	10
2002	20	1	20	1	20	1	20	1
2003	234	61	234	61	233	61	233	61
2004	72	4	72	4	59	4	72	4
2005	22	5	22	5	10	5	19	5
2006	24	10	24	10	2	1	21	10
2007	36	2	36	2	1	1	31	2
2008	57	1	57	1	4	0	44	1
2009	162	23	162	23	8	2	67	15
2010	163	89	163	89	2	2	138	58
2011	62	25	62	25	1	0	43	23
2012	46	10	46	10	4	0	30	7
2013	78	10	78	10	1	0	47	10
2014	37	2	37	2	0	0	19	1
2015	44	17	44	17	1	2	16	8
2016	106	58	106	58	2	4	16	24
2017	97	8	97	8	1	2	33	7
2018	77	10	77	10	1	0	31	5
2019	18	13	18	13	3	0	6	0

“All” refers to all unleased spells on on-shale PSF parcels since January 1, 2001. “Uncensored” spells are the subset of those parcels which were already leased on January 1, 2001. The date at which these spells begin is thus uncensored. “First” spells are the subset of all spells which are the first unleased spell for a given parcel, and “First Uncensored” spells are the subset of uncensored spells which are the first unleased, uncensored spell for a given parcel.


Xie and Liu (2005). This estimate of survival curves is similar to traditional “unadjusted” non-parametric survival curve estimate in experimental data, with the main difference being that each observation is weighted by the inverse of a propensity score. Here we estimate the probability that a spell on a parcel in a given location, of a given size, at a given point in time, is on an RAL or State parcel using a fixed-effects logit model, with Grid-by-year and Year-by-quarter fixed effects, as well as a spline in parcel acres. Letting \hat{p}_i be the estimated propensity score for spell i , we assign a weight of $\frac{1}{\hat{p}_i}$ to spells on State parcels, and a weight of $\frac{1}{1-\hat{p}_i}$ to RAL parcels. As described above, we estimate these survival curves under four different samples of unleased spells, as shown in figures A.2 and A.3.⁵² Under both 10-mile and 20-mile grid fixed effects, as well as all four sample definitions, there is no visual difference in “survival” between the two samples. Unleased spells on State and RAL parcels seem to end equally quickly.

Figure A.2: Covariate-Adjusted Kaplan-Meier Survival Curves: Grid 10 Specifications



⁵²We use the `ipw.survival` command from the R language package `RISCA`.

Figure A.3: Covariate-Adjusted Kaplan-Meier Survival Curves: Grid 20 Specifications



To formalize the test we have so far conducted visually, we compute adjusted log-rank tests of the hypothesis that two samples have the same hazard function, again informed by the ideas in [Xie and Liu \(2005\)](#). As in the adjusted survival curve estimates, the adjusted log-rank test weighs each observation by a propensity score. [Xie and Liu \(2005\)](#) shows that as long as the propensity scores are consistently estimated, the adjusted log-rank test is distributed as standard normal, allowing us to use traditional asymptotic inference techniques. Because we estimate our propensity scores using fixed effects in a high dimension non-linear model, we also compute p-values for these tests using the randomization inference technique recommended in [Xie and Liu \(2005\)](#), as well as a standard non-parametric bootstrap. Table [A.7](#) reports the test statistic values and p-values for the null hypothesis that the two survival curves are identical under each of these sample definitions, grid sizes, and inferential techniques. There is no specification in which we can reject the null hypothesis under conventional significance thresholds that State and RAL unleased spells have the same survival function.

A.4 RAL Lessor Heterogeneity

To document the extent to which different kinds of RAL lessors experience different leasing outcomes, we use lessor names to specify two notions of lessor sophistication. First, we simply count how many leases a lessor has signed in recent history, going back to the year

Table A.7: Log-Rank Tests of the Hypothesis That RAL and State Parcels Have Equally Long Unleased Spells

	All	Uncensored	First	First Uncensored
Grid 10				
statistic	-0.82	-0.72	-0.16	-0.54
p.value	0.41	0.47	0.87	0.59
ripval	0.35	0.50	0.86	0.57
bspval	0.42	0.47	0.88	0.60
N	2,892	1,408	1,788	900
Grid 20				
statistic	1.561	0.875	1.118	1.234
p.value	0.119	0.382	0.264	0.217
ripval	0.095	0.383	0.242	0.215
bspval	0.152	0.389	0.263	0.213
N	3,297.	1,579.	2,086.	1,030.

“Statistic” is the log-rank statistic for the hypothesis that RAL and State parcels have equally long unleased spells. Under the null hypothesis, this statistic is normally distributed with mean 0, and variance 1. The first p-value compares this statistic to the quantiles of the standard normal distribution. The second p-value (“ripval”) comes from a randomization inference procedure: randomly assign treatment according to the propensity scores, recompute the test statistic, and repeat 1000 times. This p-value compares the original statistic to the distribution of the simulated statistics. Finally, the third p-value, “bspval”, comes from non-parametrically bootstrapping the entire testing process 1000 times and comparing the original statistic to its standardized bootstrap distribution.

“All” refers to all unleased spells on on-shale PSF parcels since January 1, 2001. “Uncensored” spells are the subset of those parcels which were already leased on January 1, 2001. The date at which these spells begin is thus uncensored. “First” spells are the subset of all spells which are the first unleased spell for a given parcel, and “First Uncensored” spells are the subset of uncensored spells which are the first unleased, uncensored spell for a given parcel. Sample sizes differ in the top and bottom panels because the propensity score procedure will drop spells that lie in grids and/or time periods with either all RAL or all State leases, and grids of different sizes will result in a different number of dropped observations.

2000. Among all lessors who have signed an RAL lease overlying a shale formation since 2000, only 31% signed more than one lease, 5% signed five or more, and only 1% signed 10 or more. However, those highly experienced lessors signed many of the leases in our sample: 28% of RAL leases are signed by lessors who have signed 5 or more since 2000, and 13% are signed by lessors who have signed 10 or more. It seems reasonable to expect that more experienced lessors employ more sophisticated negotiation tactics. Second, we infer a lessor’s sophistication from its name. Lessors who are firms⁵³ and perhaps especially lessors that are firms explicitly working in the oil and gas business⁵⁴ may also employ more sophisticated negotiation tactics, so we generate flags for these two characteristics.

We include dummy variables for each of these categories in our standard log bonus regressions, shown in Table A.8. Across specifications, more experienced and/or more sophisticated lessors do seem to negotiate moderately higher bonus payments, but the differences are often imprecise, especially in the specifications that have 10-mile grid fixed effects. Moreover, while these lessors do better than novice negotiators, they still perform substantially worse than auctions. For example, consider specification 6, which uses 20-mile grid fixed effects and includes a dummy for lessors with 10 or more leases. These lessors negotiate bonuses that are 10 log points higher than other RAL lessors, but since the average RAL lease has a bonus that is 54 log points lower than auctioned leases, this is still 44 log points lower than than the State’s auction.

⁵³We flag a lessor as a firm if its name contains one or more of the following strings: “inc”, “ltd”, “llc”, “lp”, “corp”, “asset”, “assoc”, “company”, “holdings”, “partner”, or “manag”.

⁵⁴We flag a lessor as working in the oil and gas business if its name contains one or more of the following strings: “oil”, “gas”, “production”, “explor”, “mineral”.

Table A.8: RAL Lessor Heterogeneity and Bonus Payments

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Auction	0.54 (0.06)	0.54 (0.06)	0.54 (0.06)	0.53 (0.06)	0.54 (0.07)	0.54 (0.07)	0.54 (0.07)	0.53 (0.07)
LessorExperience5	0.03 (0.04)				0.07 (0.03)			
LessorExperience10		0.07 (0.07)				0.10 (0.05)		
LessorIsFirm			0.04 (0.04)				0.11 (0.03)	
LessorIsEP				0.00 (0.07)				0.13 (0.08)
Grid	10	10	10	10	20	20	20	20
Time	GY,Q	GY,Q	GY,Q	GY,Q	GY,Q	GY,Q	GY,Q	GY,Q
N	1,515	1,515	1,515	1,515	1,515	1,515	1,515	1,515
R^2	0.961	0.961	0.961	0.961	0.945	0.945	0.945	0.945

Regressions of the log of the bonus payment onto an auction indicator, as well as indicators for various measures of lessor sophistication. LessorExperience5 refers to lessors who lease at least 5 times during 2000-2016, and LessorExperience10 requires 10 or more leases. LessorIsFirm refers to lessors whose names suggest they are corporations, not natural people. LessorIsEP refers to lessors whose names suggest they are experienced in the oil and gas industry. All specifications include a spline in lease acres.

We also measure the extent to which the gap between auctioned and negotiated bonuses depends on the size of the lease. Our motivation for this analysis is that many of the actions involved with negotiating improved lease terms involve fixed costs which are invariant to the size of the lease. If the underlying value of a lease was exactly proportional to its size, then we might expect the owners of bigger leases to take more action to improve lease terms than the owners of smaller lessors do. Table A.9 revisits the log bonus regressions from Table 4. In column 1, we eschew our acres controls, project log bonus *per acre* onto just a dummy for auction and grid-time fixed effects. Next we include an additional dummy for whether a negotiated lease is larger than the sample median size. Columns 2 and 5 show that under this constant returns to scale assumption (ie no spline in acres), big negotiated leases do pay a modest amount more than small negotiated leases. However, there are a variety of reasons to suspect that the underlying value of a lease is not proportional to size, since modern shale drilling requires large blocks of contiguous acreage to drill long horizontal wells and to efficiently co-locate surface equipment for many neighboring wells. Columns 3 and 6 re-introduce a spline in acres, as in the main specifications in the text. The results now show that negotiated leases above the median perform no better than those below the median, relative to the bonuses achieved at auction for each size group.

Table A.9: Lease Size Heterogeneity and Bonus Payments

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.53 (0.06)	0.57 (0.07)	0.52 (0.06)	0.52 (0.07)	0.55 (0.09)	0.50 (0.08)
BigNegotiation		0.07 (0.03)	-0.02 (0.04)		0.06 (0.04)	-0.04 (0.04)
CRS Estimate	Yes G10Y	Yes G10Y	No G10Y	Yes G20Y	Yes G20Y	No G20Y
N	1,515	1,515	1,515	1,515	1,515	1,515

The variable “BigNegotiation” is equal to 1 if the lease is negotiated and if it is larger than the median negotiated lease, approximately 213 acres; otherwise it is 0. The dependent variable in the CRS specifications is the natural logarithm of bonus *per acre*, while it is the natural logarithm of the bonus in the non-CRS specifications. The non-CRS specifications in columns 3 and 6 include a spline in lease acres. All regressions include Grid10-Year or Grid20-Year fixed effects, and all include Year-Quarter fixed effects.

This analysis provides a test of the hypothesis that the negotiations in this paper perform particularly badly because of the fact that RAL lessors only receive half of the proceeds they negotiate. In this sample, the average lease above the median is nearly five times larger than the average lease below the median. This implies that the differential incentive for effort across the two groups is much larger than the differential between RAL lessors and fully private lessors. Nevertheless, even in the CRS case, the maximum amount of the auction-negotiation gap that could be explained by shirking is small. Once we flexibly control for lease size, we find no evidence that large negotiated leases perform better than smaller negotiated leases.

A.5 Lease-Parcel Comparison

In this appendix, we present comparisons of lease level and parcel level regressions run on samples which are more immediately comparable than those presented in the main text. Table A.10 presents linear results and A.11 presents Poisson results.

Columns 1-3 present lease level regressions, using 10 mile or 20 mile grid fixed effects, or DML spatial controls. Compared to the lease regressions presented in Section 5, we make two changes. First, we use the entire population of leases signed between 2004 and 2013. As described in Section 3.1, the main lease sample excludes RAL leases with features that do not overlap with auction leases, such as being too big or too small, or having multiple ownership interests. However, these lease level features are not defined at the parcel level. We thus include all leases when computing parcel level outcomes. Second, we discount all outcomes to 2004 (as opposed to the date the lease was signed). All models in Section 5 include quarter of sample fixed effects, to account for large changes in fracking technology and oil prices over time. However, the parcel regressions are inherently cross-sectional. To facilitate comparison with the parcel regressions, we eschew all time controls here.

Columns 4-6 present the corresponding parcel-level regressions. Compared to the results presented in Table 7, there is only one change. Here we restrict the sample to parcels which sign at least one lease between 2004 and 2013, since these are the parcels underlying the lease sample in the first part of the table.

Comparing across corresponding lease and parcel results reveals differences which are much smaller than the comparison discussed in Section 6.2. In levels, the bonus results are slightly higher at the parcel level, while the output results are slightly lower. In logs, the bonus results are comparable, while the lease output and revenue results are twice as large. These small differences stem from two remaining factors which are fundamentally different across the two regressions: parcel vs. lease weighting and acreage controls. All models include a spline in acres. However leases may be much smaller or much bigger than their underlying parcel(s). When they are larger, a single transaction will receive more weight in the parcel regression than other transactions on single parcels. Conversely, re-leases of a parcel after expiration will receive *less* weight in the parcel regression than in the lease regression, as will multiple concurrent smaller leases on a parcel.

Table A.10: Lease Parcel Comparison: Linear Models

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Bonus	0.354 (0.073)	0.394 (0.106)	0.352 (0.063)	0.516 (0.116)	0.559 (0.155)	0.491 (0.067)
Auction - Output	0.425 (0.159)	0.456 (0.166)	0.468 (0.186)	0.243 (0.279)	0.384 (0.257)	0.300 (0.182)
Auction - Lease Revenue	1.699 (0.597)	1.783 (0.683)	1.911 (0.712)	0.794 (1.085)	1.330 (1.049)	1.396 (0.676)
Auction - Seller Revenue	0.832 (0.178)	0.908 (0.232)	0.878 (0.199)	0.772 (0.321)	0.960 (0.339)	0.848 (0.183)
Estimate	G10	G20	DML	G10	G20	DML
Sample	All Leases	All Leases	All Leases	Leased Parcels	Leased Parcels	Leased Parcels
N	2,389	2,389	2,389	1,784	1,784	1,784

Outcomes are per acre and all models include a cubic spline in acres with 1 knot at the median lease or parcel size.

Table A.11: Lease Parcel Comparison: Poisson Models

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Bonus	0.716 (0.109)	0.777 (0.128)	0.717 (0.075)	0.690 (0.088)	0.729 (0.110)	0.645 (0.053)
Auction - Output	0.520 (0.152)	0.691 (0.153)	0.425 (0.171)	0.218 (0.154)	0.350 (0.112)	0.198 (0.114)
Auction - Lease Revenue	0.509 (0.148)	0.653 (0.168)	0.429 (0.166)	0.189 (0.140)	0.332 (0.129)	0.201 (0.112)
Auction - Seller Revenue	0.608 (0.101)	0.696 (0.116)	0.527 (0.105)	0.370 (0.087)	0.470 (0.090)	0.349 (0.063)
Estimate	G10	G20	DML	G10	G20	DML
Sample	All Leases	All Leases	All Leases	Leased Parcels	Leased Parcels	Leased Parcels
N	2,389	2,389	2,389	1,784	1,784	1,784

Outcomes are in levels and all models include a cubic spline in acres with 1 knot at the median lease or parcel size.

Appendix B Data Cleaning

B.1 Sample construction

Table A.12 presents the number of negotiated and auctioned leases existing in original data provided to us by GLO, as well as the subset that survives each of our sample restrictions. We begin with the universe of GLO leases signed between 2004 and 2016 that can be matched to original PSF parcels and have non-missing lease characteristics. We first drop leases on PSF land types other than Relinquishment Act land or belonging entirely to the GLO. Next, we use the EIA’s definition of shale formations in Texas, shown shaded in yellow in Figure A.1, to drop leases on land that does not overlie a shale formation. To ensure a comparable distribution of lease covariates, we drop especially small leases, especially large leases, and leases with especially short primary terms. We drop leases on “undivided” mineral interests, which occur when two or more parties share ownership of the surface and or mineral estate.⁵⁵ We drop a small number of leases on non-RAL land that are allocated via bilateral negotiation specifically because only one party can economically use the land. Similarly, we drop RAL leases that are allocated by auction when the State is unable to determine who the rightful surface owner is. Finally, we drop non-standard negotiated agreements between the State of Texas and E&P companies that do not have the traditional mineral lease contract structure.

Table A.12: Sample Construction

Drop Reason	Negotiation Auction	
Initial Sample	4,828	827
Not RAL or State lease	4,368	786
Not on Shale	2,517	553
Less Than 10 or Greater Than 1,000 Acres	1,985	499
Gross and Net Acreage Differ	1,658	497
Undivided Interest	1,092	493
Term Less Than 1 Year	1,073	493
Negotiated State Lease	1,066	493
Auctioned RAL Lease	1,066	454
Nonstandard Lease Category	1,061	454
Final Sample	1,061	454

Additional discussion provided in Section 3.1.

B.2 Firm Names

Though we observe the name of the firm on the lease, E&P companies sometimes use intermediaries, or “landmen,” to acquire land, and, in these cases, we might not observe the relevant firm. One reason why a firm would do this would be to prevent its competitors

⁵⁵For example, if parents John and Mary bequeath their 640 acre parcel to their two children, Bob and Jane, then Bob and Jane each have an undivided interest in the parcel. In principal, it is possible for Bob and Jane to separately lease their respective undivided interests to different oil and gas companies.

from discovering its interest in a particular play before it had had acquired enough land to develop it. This “secrecy” motivation is probably relevant, because the presence of non-E&P company lessees is much more common in the auction data than in the negotiated data. This is perhaps not surprising, since the auction records are publicly released shortly after the auction, and easily observable. To partially overcome this challenge, we use data on *lease assignments*, legal transactions which formally change ownership of a lease from one firm to another, to better infer who the ultimate E&P company is on leases initially awarded to non-E&P company lessees. We observe assignments on 18% of RAL leases and 33% of auction leases. For each non-E&P company in our data who ever assigns a lease to an E&P company, we identify a variety of “most common” assignees, using auction status, location and time. For non-E&P company leases in which we do not observe an assignment, we characterize the “real” lessee as this (conditional) most common assignee. Though this process is not perfect, it does greatly reduce the number of leases that we believe are allocated to lessees that are not E&P companies.

Appendix C Double/Debiased Machine Learning Estimation Details

To non-parametrically control for the effects of location and time on lease terms and lease outcomes, we adopt the double/debiased machine learning framework of Chernozhukov et al. (2018), which we refer to as “DML” models. In our partially linear DML models, we use the partially linear model derived in equation 4.4 of that paper. In our pseudo-Poisson models, we derive a Neyman-orthogonal score for pseudo-Poisson regression, documented below.

In all DML models, we implement the nuisance parameter estimation with random forests, using the `regression_forest` function from the `grf` package for the R language, as described in Athey et al. (2019), with 1000 trees per forest. Following the suggestions in Chernozhukov et al. (2018), we construct a single point estimate and covariance matrix of the relevant parametric terms using 5-fold cross-fitting, and report the “median” values of these across 101 randomized cross-fitted partitions, as in definition 3.3 of that paper.

To derive the Neyman-orthogonal moment for pseudo-Poisson regression, we closely follow the discussion following Lemma 2.5 of Chernozhukov et al. (2018). Let Y be a non-negative outcome, D a *binary* covariate, and X a vector of controls that we wish to model non-parametrically, like location, time and lease or parcel size. Our goal is to pick the values of θ and $\beta(X)$ that minimize a pseudo-Poisson quasi-log-likelihood criteria:

$$(\theta_0, \beta(X)_0) = \arg \max_{\theta, \beta(X)} \mathbb{E} [Y(D\theta + \beta(X)) - \exp(D\theta + \beta(X))]$$

Let $\beta_\theta(x)$ be the best fitting value of $\beta(x)$ for a given value of θ . For the above criterion, we can find this implicitly by setting the gradient of the *conditional* expectation of the criterion with respect to $\beta(x)$ equal to 0, and re-arranging terms:

$$\exp(\beta_\theta(x)) = \frac{\gamma(x)}{\exp(\theta)\delta(x) + 1 - \delta(x)}$$

where $\gamma(x) = \mathbb{E}[Y | X = x]$ and $\delta(x) = \mathbb{E}[D | X = x]$. The Neyman-orthogonal moment for pseudo-Poisson regression is the total derivative of the quasi-log-likelihood criterion with respect to θ , after we plug in the solution for $\exp(\beta_\theta(x))$:

$$\psi(Y, D, X; \theta, \gamma(x), \delta(x)) = \left(Y - \frac{\exp(D\theta)\gamma(X)}{\exp(\theta)\delta(X) + 1 - \delta(X)} \right) \left(D - \frac{\exp(\theta)\delta(X)}{\exp(\theta)\delta(X) + 1 - \delta(X)} \right).$$

This is the objective whose empirical average we set to 0 in our pseudo-Poisson regressions.

Appendix D Empirical Auction Analysis Details

We estimate the distribution of bidder values, conditional on observable and unobservable auction-level heterogeneity, in the independent private values framework⁵⁶, by following the

⁵⁶Though *offshore* mineral leasing has long been studied in the common values framework (e.g., Hendricks and Porter (1988)), most recent analysis of *onshore* mineral lease auctions takes place in the private values

model in [Roberts \(2013\)](#). The first step is to regress the natural logarithm of an auction’s reserve price R , in nominal dollars per acre, on a cubic spline in its size in acres $f(A)$, and fixed effects for the tract’s county δ_c and the date the auction took place γ_t .

$$\log R_i = f(A_i) + \delta_{c(i)} + \gamma_{t(i)} + \epsilon_i$$

We define the observed heterogeneity index as the predicted value from this regression, $H_i = \hat{f}(A_i) + \hat{\delta}_{c(i)} + \hat{\gamma}_{t(i)}$, and the unobserved heterogeneity index as $U_i = \log R_i - H_i$, the residual. We estimate this model on the entire population of GLO auctions for tracts in counties that overlap shale formations, during 2004-2016, including auctions that do not receive any bids at or above the reserve price.

Next, we use (H_i, U_i) as the heterogeneity covariates in the procedure specified in section 4 of [Guerre et al. \(2000\)](#). After recovering the bidder values that are consistent with equilibrium bidding, we find that the bid-value relationship is occasionally non-monotonic. When this occurs, we follow the advice in [Chernozhukov et al. \(2009\)](#) and re-order estimated values within an auction to guarantee monotonicity.

Table [A.13](#) shows how our key auction statistics are sensitive to different kernel bandwidth choices, dropping bids near the boundary of the support of the distribution of bids and/or heterogeneity covariates, and dropping auctions where value re-arrangement is necessarily to maintain monotonicity of the bid-value relationship.

Table A.13: Auction Robustness

Sample	Bandwidth Scaling	Auctions	Average 1-to-2 Markup	Average 1-to-2 Allocative Gain
All	0.5	202	0.34	0.41
All	1.0	202	0.34	0.47
All	2.0	202	0.34	0.58
Drop Trimmed	1.0	154	0.32	0.40
Drop Rearranged	1.0	162	0.36	0.49
Drop Trimmed and Rearranged	1.0	117	0.34	0.41

This table summarizes bids from GLO auctions with 2 or more bidders across various assumptions made in the GPV analysis. The sample includes all auctions in the first three rows. In the fourth row we drop auctions where one or more bids would be trimmed under the GPV trimming rules. In the fifth row we drop auctions where the estimated pseudovalues are not monotonic in the bids. In the sixth row, we drop auctions where bids would be either trimmed or pseudovalues would need to be rearranged to impose monotonicity. “Bandwidth Scaling” refers to whether we use the Silverman rule for bandwidth selection (1.0) or half (0.5) or twice (2.0) the Silverman rule. The “1-to-2 Markup” is the natural logarithm of the ratio of the winning bid to the second highest bid. The “1-to-2 Allocative Gain” is the natural logarithm of the ratio of the GPV pseudovalues of the winning and second highest bids.

framework. [Kong \(2021\)](#) and [Kong \(2020\)](#) both provide informative discussions about why private values make sense in studying the recent shale boom.

Appendix E RAL Lease Addenda

In addition to specifying bonus payments, royalty rates and primary terms, mineral leases also specify how the contracting parties will resolve disagreements about issues related to environmental impact, on-site water usage, and surface property disruptions, among other things. These protective clauses are standardized in the GLO auction lease agreement, and there are “default” values for them in the GLO’s required RAL lease agreement. However, RAL surface owners and their contracting partners can optionally negotiate some deviations from the standard lease. To the extent that RAL surface owners are willing to forego up-front bonus payments for stricter surface protections during subsequent exploration and production, we might be worried that the differences in bonus payments that we observe are not caused by the mechanism itself, but rather by a compensating differentials story.

It is important to note that it is not possible to directly condition on the content of these protective clauses and maintain our causal interpretation of the differences between auctioned and negotiated leases, for the same reason that we do not condition on royalty rates and primary terms. See the discussion at the end of Section 5. However, having recognized the challenge in interpreting these results as causal, we explore the extent to which RAL lease addenda content affects our regression estimates.

To do this, we had a team of research assistants do a dual-entry review of the text of these lease addenda for all RAL leases signed between 2005 and 2016. They characterized the extent to which each one improved or deteriorated the surface owner’s rights along dimensions such as environmental impact, water usage, and surface property disruptions, guided by the procedure in [Vissing \(2017\)](#). About 73% of RAL leases have one or more additional clauses in their lease addenda. In [Table A.14](#), we include measures of these protective clauses in bonus and output regressions like those shown in [Table 4](#) and [6](#). The first two columns mirror the result shown in the main text: auctioned leases pay 50 or more log points in up-front bonus payments than negotiated leases do, and output is similarly higher. In the next two columns, we include covariates which measure the number of pages in an RAL lease’s addendum, as well as the number of specific legal clauses documented. Finally, in the last two columns we include covariates for each specific kind of clause that occur in these addenda, coded as -1 if a lease’s addenda deteriorates the surface owner’s rights, relative to the standard RAL lease, 0 if it is absent or does not affect the surface owner’s rights, and $+1$ if it improves upon the surface owner’s rights. Across all specifications, we find no evidence that variation between auctioned and negotiated leases in protective clauses can “explain away” the observed differences in both bonus payments and output.

Table A.14: Bonus Payments, Lease Output, and Mechanism Type: Robustness to RAL Lease Addenda

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - log(Bonus)	0.57 (0.06)	0.65 (0.05)	0.61 (0.06)	0.80 (0.07)	0.61 (0.06)	0.70 (0.06)
N	1,409	1,409	1,409	1,409	1,409	1,409
Auction - Output	1.10 (0.34)	1.39 (0.44)	1.23 (0.42)	1.66 (0.52)	1.16 (0.43)	1.32 (0.44)
N	1,153	1,153	1,153	1,153	1,153	1,153
Addenda Controls	None	None	Pages + Clauses	Pages + Clauses	Individual	Individual
Grid	10	DML	10	DML	10	DML
Time	GY,Q	DML	GY,Q	DML	GY,Q	DML

Columns 1 and 2 report baseline specifications with no controls RAL lease addenda content. We measure lease addenda content in columns 3 and 4 using the number of pages and number of individual clauses in the addenda. In Columns 5 and 6, we include a set of individual dummy variables indicating whether the addenda improves (+1), deteriorates (-1) or is similar to the standard auction lease (0) along each of these margins: Surface Protection, Payment Terms, Location Requirements, Pugh Clause, Cleanup Terms, Livestock Protection, On-site Water Use, Waste Management, Definitional Changes, Pollution Protection, Infrastructure Constraints, Caliche Use, Additional Fees, Time Constraints, Miscellaneous. Columns 1, 3, and 5 control for space and time using 10-mile grid by year of sample fixed effects, as well as fixed effects for quarter of sample. Columns 2, 4, and 6 use a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#), with non-parametric controls for lease latitude, longitude, effective date, size and addenda measures. Sample includes leases signed between 2005-2016.