

NBER WORKING PAPER SERIES

RELINQUISHING RICHES:
AUCTIONS VS INFORMAL NEGOTIATIONS IN TEXAS OIL AND GAS LEASING

Thomas R. Covert
Richard L. Sweeney

Working Paper 25712
<http://www.nber.org/papers/w25712>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2019, Revised October 2019

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 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Thomas R. Covert and Richard L. Sweeney. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Relinquishing Riches: Auctions vs Informal Negotiations in Texas Oil and Gas Leasing
Thomas R. Covert and Richard L. Sweeney
NBER Working Paper No. 25712
March 2019, Revised October 2019
JEL No. D44,L13,Q35

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 40 percent larger upfront payments and 60 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.

Thomas R. Covert
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
thomas.covert@chicagobooth.edu

Richard L. Sweeney
Economics Department
Boston College
Chestnut Hill, MA 02467
sweeneri@bc.edu

Code for replication is available at https://github.com/rlsweeney/public_cs_texas

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 almost thirteen hundred leases for publicly owned minerals in Texas between 2005 and 2016, we find that auctioned leases sell for 36 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 auctioned and negotiated parcels have a successful transaction. The economic significance

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.

of these results is large: on average, negotiated leases would have generated \$133,000 more up front revenue had they transacted using an auction.

In principle, this improvement in up front 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 46 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* seller revenue by about \$351,000 per lease. Moreover, we show that while auctions allocate minerals to different firms, both the payment and output differences hold *within* firm, suggesting that firm-lease “match” plays an important role in determining lease outcomes, and that auctions generate better matches than negotiations do. As we discuss in Section 7, the fact that the optimal partner varies from lease to lease implies that the gains from attracting additional bidders and from using a formal mechanism to select between them are large in this setting.

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 \(2017\)](#) 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 by directly comparing observable measures of these outcomes, 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

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.

costly landowner search effort (Hortaçsu and Syverson, 2004; Allen et al., Forthcoming; Cuesta and Sepulveda, 2018), bilateral bargaining (Backus et al., 2015, 2018; Larsen, 2018) and even take-it-or-leave-it behavior on the part of some buyers. 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.³ 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 (2018). 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 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, 2017), measure bidder uncertainty about competition (Kong, 2016), and evaluate the choice of “security” sold by the winning bidder to the auctioneer (Bhattacharya et al., 2018). 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., 2017). In our setting, we directly observe the full set of payments received by mineral owners, including bonus payments and royalty rev-

³Indeed, our results are consistent with a variety of theoretical comparisons between simple auctions and more complicated alternatives. We explore these interpretations in Section 7.

enue. We find that bonus payments represent more than seventy percent of total landowner revenue earned to-date for the average lease, and by construction, they are the entirety of landowner revenues for the two thirds 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, and Sections 5 and 6 present the results. 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 unique 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.⁴ 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 land owners and E&P companies are formalized through mineral lease agreements, which are contracts with three key elements: a *primary term* before which the lessee (the E&P company) must begin drilling; a *royalty rate* providing the lessor (the landowner) with a share of any realized drilling revenues; and an upfront *bonus payment* to secure the right to explore.⁵ Many leases also include *delay rentals*, which are payments the

⁴According to EIA data, oil prices were \$66.73 per barrel (Brent) and natural gas prices were \$3.69 per million BTU (Henry Hub).

⁵This 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. (2018).

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).⁶

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 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.

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, including more than 100 that drilled at least 1 well per month. Compared to other capital intensive businesses, the oil and gas industry in America is not concentrated, with the 50 largest firms drilling about half of all wells drilled in Texas. The largest firm, XTO Energy, drilled fewer than 5% of all new wells.⁷ 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.

Many sources of heterogeneity among E&P companies are "vertical" in nature, in that some firms have either consistently lower costs or higher productivity than others. There are wide differences in firm size and observable measures of firm sophistication among the set of active firms in the US onshore E&P business. Indeed, some of the largest companies in the world, like Exxon and Chevron, compete for leases against thousands of privately held E&P companies with fewer than 500 employees. Beyond observable differences in firm

⁶We study these "lease addenda" formally in Appendix C.

⁷XTO Energy is a subsidiary of Exxon Mobil, acquired in 2010.

size and sophistication, there is heterogeneity across E&P companies in their decisions to hire external service contractors to perform drilling and completion services or to maintain these capabilities in house. 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). Finally, it is possible that some firms may simply be able to process post-acquisition lease information more effectively, and in doing so, more efficiently select which of their leases to drill.

In addition to these vertical differences in E&P company quality, there are also many potential sources of *horizontal* heterogeneity across firms, which may make some better at developing a particular piece of land than others. For example, firms who already control acreage in one area may be able to develop drilling plans that minimize the number of wells necessary to extract minerals, relative to firms who have less existing nearby acreage holdings. Firms who own hydrocarbon transportation infrastructure close to a given parcel may experience cost advantages in developing that specific parcel, but not other parcels further away from this infrastructure. 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.

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 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.⁸ To stave off “armed rebellion” by these

⁸Although 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

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 1931.⁹ 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.

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.

The State 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 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 the Texas General Land Office (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, so the cash bids are analogous to the bonus payment on a negotiated lease.

3 Data

From the Texas General Land Office, we obtained data on the universe of oil and gas leases signed between 2005 and 2016 on PSF lands where the State of Texas retained full or partial mineral rights. This includes leases on RAL parcels, sold by the State prior to 1931, and

advent of the oil age in the United States (Yergin, 2008).

⁹The Texas Supreme court finalized their interpretation of the Relinquishment Act on June 29, 1931.

parcels that were sold after 1973 or still belong to the State.¹⁰ Our initial dataset includes the shape, location, size, effective date, bonus payment, primary term and royalty rate for 4,071 RAL leases and 646 State auction leases. For all leases that eventually result in drilling, we observe monthly payments for gas and oil royalties remitted to the state, up through March, 2019. We use this royalty payment data to both measure the output that a lease generates, and also determine which leases were drilled, without having to spatially match drilling data to leases. We combine the royalty revenue data with output price¹¹ information to infer monthly oil and gas production for drilled leases.

We spatially intersect this lease-level dataset with a parcel map of all lands in the PSF. We acquired this map from P2Energy Solutions, a private contractor hired by the State of Texas 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 5.2.

The GLO uses first price, sealed bid auctions to allocate its non-RAL leases. GLO 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. 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.¹² We then exclude leases smaller than 10 acres or bigger than 1,000 acres, as parcels in the original Texas Permanent School Fund 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 bequeathment subsequent to privatization. Finally we exclude a small number of leases with missing or inconsistent information. We provide additional information on these restrictions, as well

¹⁰The State did not retain primary mineral ownership on Free Royalty lands sold between 1931 and 1973. As a result, the state does not track full payment information on these parcels.

¹¹We use the Henry Hub price for natural gas royalties and the WTI price for oil royalties.

¹²We use the EIA’s definition of shale formations in Texas, shown shaded in yellow in Figure A.1.


as an associated waterfall table with the number of leases surviving each cut, in Appendix B.1. The resulting dataset of 846 negotiated leases and 428 auctioned leases is summarized in Table 1. Figure 1 demonstrates the distribution of lease types over time.

Table 1: Lease Summary Statistics by Type

Variable	Negotiation (N = 846)				Auction (N = 428)				Difference	p-value
	mean	sd	min	max	mean	sd	min	max		
Land Characteristics										
Acres	0.28	0.26	0.01	0.99	0.36	0.24	0.01	0.73	-0.08	0.00
Shape Quality	0.94	0.14	0.10	1.00	0.96	0.10	0.28	1.00	-0.01	0.05
Multiple Parcels	0.05	0.22	0.00	1.00	0.03	0.16	0.00	1.00	0.03	0.02
Shale Thickness	3.18	1.27	0.15	5.98	3.79	1.21	0.10	5.84	-0.60	0.00
Lease Characteristics										
Bonus	1.10	1.19	0.04	10.41	2.19	2.64	0.02	15.40	-1.09	0.00
Term	3.85	1.16	1.00	5.00	4.75	0.66	3.00	5.00	-0.90	0.00
Contracted Rentals	0.22	0.47	0.00	2.82	0.32	0.54	0.02	2.27	-0.10	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										
Drilled	0.38	0.49	0.00	1.00	0.33	0.47	0.00	1.00	0.06	0.05
Output	0.22	0.53	0.00	4.67	0.23	0.51	0.00	3.17	-0.01	0.72
Lease Production Revenue	7.93	18.95	0.00	153.49	8.76	20.61	0.00	169.18	-0.83	0.49
Realized Rentals	0.08	0.21	0.00	2.25	0.15	0.33	0.00	2.27	-0.07	0.00
Total Seller Revenue	3.13	5.03	0.04	43.37	4.49	5.83	0.03	44.44	-1.37	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 (176 negotiation leases and 28 auction leases); bonus, lease revenue, and total seller revenues are all reported in thousands of nominal dollars per acre; output is reported in thousands of barrels of oil equivalent per acre; term is reported in years; contracted delay rentals and realized delay rentals are reported in thousands of dollars per acre. *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, realized delay rentals, and the mineral estate’s royalty share of lease revenues.

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 these measures are small. They also generate substantially higher bonus payments (per acre) and pay slightly higher royalty rates and delay rentals, while auctioned leases have longer primary terms. Auctions are slightly less likely to be drilled, produce equivalent amounts of output and associated lease revenue, and the difference in total revenues (the sum of bonus payments, royalty income on production, and realized delay rentals) is slightly larger than the difference in bonus payments. Figure 1 shows that auctions are not consistently prevalent over time. In particular, there are relatively few auctions in 2009 (when oil prices temporarily crashed during the financial crisis) and in 2012 (when gas prices reached lows not seen in a decade). Appendix Figure A.1 shows that auctioned and negotiated leases are also not evenly distributed across space, except possibly in West Texas, where the Permian Basin shale play has recently experienced a surge in leasing activity. These differences in timing and location underscore the importance of flexibly controlling for these factors in our empirical specifications below.

4 Empirical Strategy

We use the variation in leasing mechanisms employed on parcels initially placed in the Permanent School Fund in 1895 to measure how auctions affected lease outcomes, relative to informal negotiations, during the recent fracking boon. In the ideal experiment, we would have randomized mechanism type, auction or informal negotiation, among a population of

private mineral owners on top of shale formations, on the eve of the fracking boom. In practice, our sample consists of leases signed between 2005 and 2016 on PSF lands where the State retained a mineral interest. Within this sample, mechanism type is determined not by randomization, but by the date on which the parcel underlying each lease was first sold by the State. As we argue below, variation in privatization dates is unlikely to be correlated with unobservable determinants of lease outcomes in the modern shale era.

In addition to the fact that leases are quasi-experimentally assigned between the two mechanisms, our comparison of auctioned and negotiated leases in PSF lands is appealing because the two types of leases are similar on all other important contracting dimensions. Both types of leases are buyer initiated, with an E&P company approaching a landowner or nominating a parcel to the State. The language in the two types of contracts is extremely similar because the State requires RAL surface owners to use a standard lease document, whose structure is nearly identical to the lease contract the State uses in its auctions.¹³ Finally, on RAL leases, the State of Texas, rather than the private surface owner, represents the mineral estate in legal matters against lessees. Thus, contractual disputes that occur after a firm signs a lease are between that firm and the State, regardless of which leasing mechanism is used. For these reasons, the property right that E&P companies buy in RAL leases is effectively identical to what they buy in auctioned leases.

With that background in mind, we estimate several versions of the following 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 contract details. In some specifications, we also condition on 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 the two primary determinants of lease outcomes: *where* leases are and *when* they transact, which we write as $\delta_{L,T}$ in equation 1. Leases on parcels with better mineral resources may transact at higher prices, attract more investment and produce more output. Similarly, leases that transact during periods of high output prices or increased technological progress may earn higher prices or generate better post-leasing outcomes. To ensure that differences in resource quality across space do not

¹³This is in stark contrast to the broader private mineral leasing market, where contractual terms vary considerably (Timmins and Vissing, 2017). 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.

confound our comparisons, we include fixed effects for 10 mile by 10 mile square “grids” that contain a lease’s centroid. Figure 2 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. To account for time varying unobservable determinants of lease outcomes, we also include fixed effects for the year-quarter of a lease’s transaction. In some specifications, we restrict comparisons to leases in the same grid and year-quarter.

Figure 2: 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 location and time on lease outcomes, and, as Figure 2 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).¹⁴ Both of these strategies ensure we are making comparisons between leases with

¹⁴Specifically, we estimate a Robinson (1988) style partially linear specification of equation 1. We follow

similar mineral quality that transact at similar times. To interpret estimates of τ in equation 1 as the causal effect of auctioning vs. negotiating a lease, we must thus assume that the allocation mechanism is independent of any *residual* determinants of lease outcomes, those that remain after conditioning on location, time, and other observable characteristics.

To understand this assumption, recall that the transaction mechanism for a lease is determined by the time at which the State of Texas sold the PSF parcel beneath it. Figure 3 shows a timeline of this process. 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 those leases), are both uncorrelated with residual determinants of value in the shale era.

Figure 3: 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 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 $\theta = (\tau, \beta)$ that minimizes the “cross-fitted” empirical analogue of

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

where $D = (\text{Auction}, X)$, $\gamma(l, t) = \mathbb{E}[Y | L = l, T = t]$, $\delta(l, t) = \mathbb{E}[D | L = l, T = t]$ and we estimate the functions $\gamma(\cdot)$ and $\delta(\cdot)$ using random forests. For more details, see Appendix A.7.

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 the rock is and the concentration of hydrocarbons within it. However, the assignment process is complete by 1973, decades before 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 leases 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 3 is similar on observable measures of quality. Table 2 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.

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 2 estimate similar regressions, using parcel size, shape and distance to water as the outcome variable. Both

Table 2: Parcel comparison: Land in the PSF Overlying Shale Formations

	Thickness	Acres	Shape	Water	Rivers
Auction	0.013 (0.055)	-0.087 (0.016)	-0.011 (0.006)	0.184 (0.413)	0.051 (0.073)
Free Royalty	-0.013 (0.041)	-0.077 (0.016)	0.000 (0.005)	-0.177 (0.374)	0.001 (0.057)
Average	3.147	0.296	0.960	12.089	0.621
N	2,461	3,488	3,488	3,488	3,488
R^2	0.895	0.441	0.416	0.812	0.543

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.

auction and Free Royalty parcels are smaller than RAL parcels. Because of this, all of our lease level regressions flexibly control for lease acreage. Auctioned parcels are also somewhat less convex than Free Royalty and RAL parcels, but the difference is small relative the average PSF parcel shape and unlikely to be economically significant. Finally, all three parcel types are similarly close to standing sources of water and to rivers and streams.

5 Seller Revenue Results

We begin by investigating the impact of auctions on seller revenues, estimating several versions of equation 1 with the natural logarithm of bonus payments as the dependent variable. Table 3 presents the results. All models include controls for primary term, royalty rate and acres.¹⁵ 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 36 log points more in bonus payments than similar negotiated leases, and this difference is precisely estimated.¹⁶ In column 2, we interact the

¹⁵Note that we do not condition on contracted delay rental payments here because they are mechanically correlated with bonus payments for both types of leases. Instead, we include realized delay rental payments as a component of total seller revenues in Section 6. In Appendix A.5 we also formally measure the differences in contracted and realized delay rentals using the framework in equation 1.

¹⁶Note that in percentage terms, this difference considerably larger than 36%, as $\exp(0.36) - 1 \approx 43\%$. In appendix A.3, we repeat these regressions in levels, with the dollar value of the bonus payments (per acre) as the left-hand side variable.

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 36 log points more, and is still precisely estimated. This model, which compares leases for minerals that are located at roughly the same place and which transact at roughly the same point in time, is our main specification.

Table 3: Bonus Payments and Mechanism Type

	(1)	(2)	(3)	(4)	(5)
Auction	0.36 (0.07)	0.36 (0.08)	0.37 (0.12)	0.37 (0.08)	0.52 (0.05)
Grid	10	10	10	20	DML
Time	Q	GY,Q	GYQ	GY,Q	DML
N	1,274	1,274	1,274	1,274	1,274
R^2	0.866	0.953	0.973	0.917	

The dependent variable in each regression is the natural logarithm of the 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. \(2018\)](#). All models include a spline in acres and linear terms for term length and royalty rate. The average negotiated bonus payment is \$1,096 per acre.

In the remaining 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. Finally, 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 \$133,000 larger up-front bonus payment.

5.1 Robustness

As we discussed in Section 4, our key identifying assumption is that land that was initially owned by the state but sold by 1931 is similarly valuable for today’s hydrocarbon exploration as land from the same allocation that was not sold as of 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 2 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 85 years 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 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 convex hull containing it and determine whether the lease spans more than one distinct parcel. Next, we measure the distance of each lease to road infrastructure using GIS data from the Texas Department of Transportation. Finally, we compute the most common surface coverage characteristics of a lease using the National Land Cover Database.¹⁷ In Table 4, we confirm that the bonus results above are robust to including these characteristics as controls at the lease level. Model 1 repeats the grid-year fixed effect model from Table 3, 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. Column 2 estimates the same specification using the DML method to control for location and time. Columns 3 and 4 repeat these specifications, but include the thickness of the shale underlying the parcel as well. Across all of these specifications, we continue to find that auctions pay significantly more than negotiations do.¹⁸

¹⁷Appendix Table A.1 demonstrates that leases the parcels underlying our sample leases are statistically indistinguishable along these surface measures.

¹⁸We also estimate overlap-weighted treatment effects in Appendix A.4.

Table 4: Bonus Payments and Mechanism Type: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.37 (0.08)	0.51 (0.05)	0.37 (0.08)	0.45 (0.05)	0.47 (0.11)	0.62 (0.07)
Estimate	G10Y	DML	G10Y	DML	G10Y	DML
Surface Controls	Yes	Yes	Yes	Yes	No	No
Thickness Controls	No	No	Yes	Yes	No	No
Private Only	No	No	No	No	Yes	Yes
N	1,274	1,274	1,070	1,070	1,073	1,073
R^2	0.954		0.954		0.956	

The dependent variable in each regression is the natural logarithm of the bonus payment. Surface controls include shape regularity, a dummy variable for whether the lease spans multiple parcels, surface cover measures, and distance to roads and water sources. In columns 3 and 4, the sample excludes leases overlying parts of the Eagle Ford Shale and leases overlying the the Haynesville and Barnett shales, for which there is no thickness information available. In columns 5 and 6, the sample is restricted to leases with private surface ownership. Columns 1, 3, and 5 use fixed effects for year-by-10-mile grid, as well as quarter of sample, with standard errors clustered by grid. Columns 2, 4 and 6 use a double/debiased Machine Learning routine to control for location and time, as recommended in [Chernozhukov et al. \(2018\)](#). All models include a spline in acres, and linear terms in term length and royalty rate.

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 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. If this were the case, it 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 5 and 6 of Table 4 present our main bonus regressions specifications re-run 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 Table 3: both point estimates are statistically significant at conventional levels, and they comfortably lie within the confidence intervals for the estimates in columns 2 and 5 of Table 3. 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 C, 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.

5.2 Extensive Margin Considerations

The results in Table 3 show that auctioned transactions occur at substantially higher prices than negotiated transactions. However, this is a comparison between *successful* transactions, and not all *attempted* transactions are successful: auctions fail if they attract no bids at or above the posted reserve price, and negotiations analogously fail when surface owners demand lease terms that exceed the willingness to pay of their contracting partners. When attempted transactions fail, the short-run welfare of landowners and their potential contracting partners is effectively zero. If failures are common, and differentially likely across the two mechanisms, the true seller gains from auctions could be quite different from the observable revenue difference conditional on leasing. Thus, to correctly interpret our revenue differences, we must first check for the presence of differences across mechanisms in the probability of a successful transaction.

For auctioned leases, we can directly compute the probability of a successful transaction, because we observe the list of parcels that go up for auction, as well as the subsequent bids. Among GLO auctions on PSF land, 44% of nominated parcels failed to receive a qualifying bid. So on a per potential transaction basis, failure is quite common. The GLO often offers

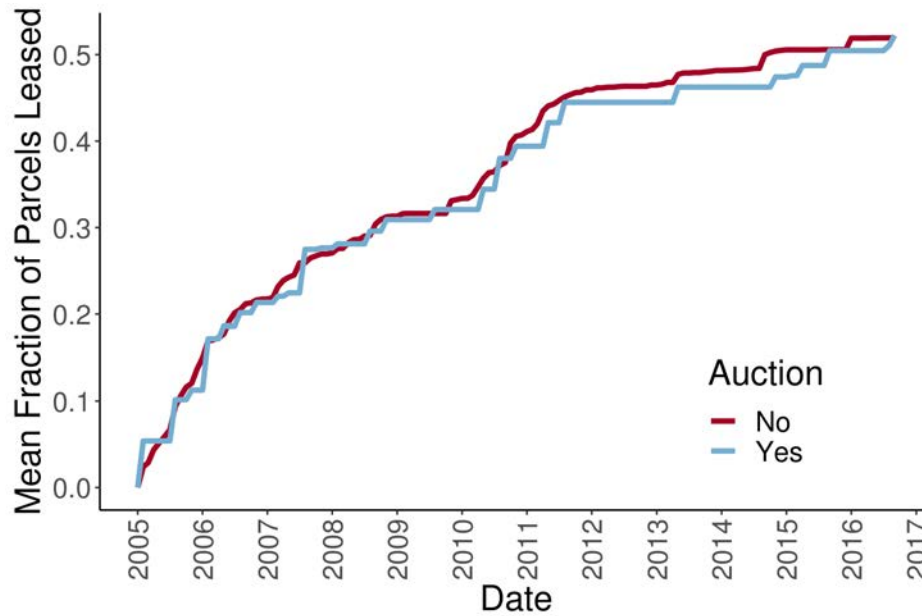
to sell these failed parcels again in future auctions, to the point that 73% of all observed nominated parcels transact *at some point* in our sample. Given that auctions don't always clear, even after repeated attempts at transaction, it could be the case that the difference in seller revenues we observe on successful transactions could be offset by a higher likelihood of transaction among RAL negotiations.

Unlike auctions, we don't observe attempted RAL negotiations that fail, so we observe neither the likelihood of "nomination" nor the probability of successful transaction, conditional on being nominated. However, we can still characterize the total extensive margin differences between auctions and negotiations, inclusive of both differences in nomination and transaction success, by comparing the number of parcels that could ever have a transaction under a given mechanism, with the number of those parcels on which we actually observe a lease.

We first visualize the rate at which auction and negotiation parcels are leased over time in Figure 4. Using the same sample as Table A.1, we compute the fraction of auction and negotiation parcels within each 10-mile grid that have been leased at least once by the start of a given month, and plot the average across grids for each month between January 2005 and December 2016. Visually, the arrival rate of a parcel's first successful transaction is comparable across the two mechanisms, providing initial evidence that differences in the nomination process or probability of a successful transaction are unlikely to be important.

To ensure that differences across parcels in size, shape quality, land cover characteristics, or distance to infrastructure don't mask differences in the likelihood of a successful lease, we also report estimates of parcel-level regressions in Table A.6. The left-hand-side variable is a dummy indicating that at least one lease occurs during our sample period (2005-2016). In the first four columns, which compare auctioned and negotiated parcels in our entire sample using both fixed-effect and DML methods, we find no evidence that parcels which lease by negotiation are any more likely to lease than parcels which lease by auction, consistent with Figure 4. In columns 5 and 6, we restrict to the sample of parcels over which we have shale thickness information, and in this subsample, there is a small but precisely estimated difference in the likelihood of leasing, but it favors auctions, not negotiations. Given these results, it does not appear that the larger seller revenues we observe in successful auctions are offset by a lower likelihood of transaction. If anything, parcels governed by auctions are somewhat more likely to transact.

Figure 4: Time to First Lease for Auction and RAL Parcels



Average across 10 square mile grids of the fraction of parcels that have leased at least once since January 2005, by parcel type.

Table 5: Likelihood of Leasing and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.008 (0.019)	0.007 (0.015)	0.010 (0.018)	0.009 (0.016)	0.033 (0.017)	0.031 (0.015)
Grid	10	DML	10	DML	10	DML
Surface Controls	No	No	Yes	Yes	Yes	Yes
Thickness Controls	No	No	No	No	Yes	Yes
N	1,763	1,763	1,763	1,763	1,202	1,202
R^2	0.810		0.813		0.537	

The dependent variable equals 1 if a parcel was ever leased and 0 otherwise. Surface controls include shape regularity, a dummy variable for whether the lease spans multiple parcels, surface cover measures and distance to roads and water sources. In columns 5 and 6, the sample excludes parcels overlying parts of the Eagle Ford shale and parcels overlying the Haynesville and Barnett shales, for which there is no thickness information available. Columns 1, 3 and 5 use fixed effects for 10-mile grids with standard errors clustered by grid, while columns 2, 4 and 6 use a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). All models include a spline in the size of the parcel in acres.

6 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.¹⁹ Alternatively, auctions and negotiations might select equally productive firms, but the absence of formality in the negotiation process allows negotiation winners to pay less than they would in an auction. Under this explanation, 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. 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 the 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. Though our data covers all production between January, 2005 and March, 2019, leases are expected to produce output for 20 years or more, so the discounted sum of realized revenue is right censored, even for the earliest cohort of leases. To limit the bias generated by this censoring, we focus on leases whose primary term has concluded by the end of our royalty data, so that all leases are at least properly categorized as having drilling or not. 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 similarly censored.

The first row of panel (a) in Table 6 presents these results. The model specifications in each column are identical to those in Table 3, showing the effects of mechanism type on each output measure, under various spatial, temporal, and surface characteristic controls. Across these specifications, auctioned leases produce \$3,500 to \$6,100 more in discounted lease production 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 discounted production revenues from negotiated leases of \$8,308 per acre.

¹⁹Initial misallocation may not be remedied through subsequent reassignment due to asymmetric information (Brehm and Lewis, 2019).

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

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Lease Revenue	4.54 (1.81)	3.52 (2.05)	4.07 (3.81)	3.85 (1.82)	6.13 (1.84)	3.68 (2.13)
R^2	0.524	0.766	0.853	0.612		0.749
Auction - Output	0.12 (0.04)	0.11 (0.05)	0.12 (0.10)	0.12 (0.04)	0.15 (0.04)	0.11 (0.05)
R^2	0.512	0.752	0.833	0.598		0.711
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	1,175	1,175	1,175	1,175	1,175	974

(b) Pseudo-Poisson Regression Models

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Lease Revenue	0.44 (0.16)	0.46 (0.18)	0.48 (0.26)	0.36 (0.20)	0.60 (0.19)	0.46 (0.16)
Auction - Output	0.48 (0.15)	0.55 (0.19)	0.56 (0.27)	0.46 (0.22)	0.60 (0.17)	0.54 (0.18)
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	1,093	738	613	944	1,175	618

The dependent variables are the discounted sum of oil and gas production revenue (Lease Revenue) and discounted barrels of oil equivalent (Output), both in thousands. In the top panel, the estimates come from linear regression models, while in the bottom panel, they come from pseudo-poisson quasi-maximum likelihood models. 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. \(2018\)](#). All models except for column 5 in the bottom panel include a spline in acres, and linear terms in term and royalty rate. The DML models in the bottom panel include acres, term and royalty rate as covariates in the random forest routines. “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 the thickness of the shale formation. 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 \$8,308 in lease revenue per acre, and 233 of discounted barrels of oil equivalent 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.²⁰ 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.²¹ 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 110 to 150 additional discounted barrels of oil equivalent, per acre. These estimates are also more precise than the revenue estimates, because they do not include the variation coming from 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 more than 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, as described above, 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).²² 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,²³ we also estimate pseudo-poisson regression models, which effectively project the logarithm of the expected value of revenue (or output) onto mechanism type, covariates, and location and time:

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

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

²⁰For a review of this literature, see Syverson (2011).

²¹We assume 1 thousand 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>

²²Appendix Table A.7 estimates the effects of mechanism type on the likelihood that a lease is drilled. The results are noisy, but suggest large increase on this extensive margin.

²³Log 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.”

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 36 to 60 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.

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 Table 7, which shows regressions of the sum of bonus payments, realized delay rentals and discounted royalty revenues onto the same right hand side variables as in Table 6. Auctions generate \$1,200 to \$2,200 more revenue per acre for landowners than the negotiations do. Under our main specification (column 2), this difference is worth about \$351,000 for the average RAL lease. These results show that auctions have an economically enormous impact on sellers.

Table 7: Total Seller Revenue and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	1.49 (0.47)	1.24 (0.57)	1.53 (1.04)	1.39 (0.49)	2.25 (0.47)	1.27 (0.58)
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	1,175	1,175	1,175	1,175	1,175	974
R^2	0.568	0.795	0.873	0.653		0.780

The dependent variable is the discounted present value of the sum of bonus payments, delay rentals paid, and production royalties, per acre. In columns 1-4 and 6, 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. Column 5 uses a double/debiased machine learning routine, as recommended in Chernozhukov et al. (2018). For estimation details, see A.7. All models include a spline in acres, and linear terms in term length and royalty rate. “Extra” controls include shape regularity, a dummy variable for whether the lease spans multiple parcels, surface cover measures, and distance to roads and water sources. The sample includes all leases whose primary term ends before March, 2019.

6.1 Unpacking the Productivity Results

Table 6 provides evidence that auctions allocate leases to users who are more likely to drill them, and who produce more output with them. In this section, we provide statistical evidence regarding the relative contribution of vertical or horizontal productivity differences between firms in generating these results. The motivation for this decomposition is to understand whether or not auctions are *necessary* to deliver the gains we estimate above. If the gains from auctions come primarily from identifying persistently productive *firms* (in the vertical sense), one can imagine relatively light-handed policy interventions which would steer landowners towards better firms (or away from worse firms). However, if the latent firm productivity ordering varies from parcel to parcel, it is difficult to imagine a way to achieve higher allocative efficiency without widespread adoption of some auction-like mechanism.

As we discussed in Section 2, there are a number of reasons to suspect that some E&P companies are persistently more productive than others. Since the identity of the winning firms is easily observable, we first check whether auctions and negotiations pick *different* firms.²⁴ We do this by tabulating auction and negotiation “market shares” for each of the ten most active lessees, as shown in Table 8.²⁵ For these especially active lessees, a firm’s share of leases in the auction market is quite different than its share in the negotiation market. The data soundly reject a Chi-squared test of the hypothesis that a firm’s auction market share is the same as its negotiation market share ($p < 2 \times 10^{-16}$).²⁶

The net result of these differences is differential concentration in lease ownership across the two mechanisms. Table 8 suggests that the auction market is more concentrated than the negotiation market: the top 10 auction winners won 56% of all auctions, while the top 10 negotiators won just 45% of all negotiations. If these firms at the top are also persistently more productive, then it would be consistent with the idea that auctions generate more output because they select better *firms*.

Having established that auctions allocate to different firms, we next ask whether vertical differences between firms can explain our results. If auctions produce more because the firms that win auctions produce more on both types of leases, then comparisons of output

²⁴Note that the comparison here is the frequency with which firms win, not whether some firms only participate in auctions and others only use negotiations, as in [Bajari et al. \(2009\)](#). Auctions are a small fraction of the broader minerals market in Texas, which is all negotiated and in which all auction participants are also active.

²⁵Firm identities in our data are recorded with some error (typos, etc). We describe our process for cleaning these names in [Appendix B](#).

²⁶Chi-squared 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: Top 10 Auction Winners and Negotiators

Firm	Leases	Auction Share	Negotiation Share
CHESAPEAKE	112	0.192	0.035
ENERGEN	78	0.065	0.059
LEWIS OPERATING	73	0.005	0.084
PETROHAWK	71	0.091	0.038
PETRO HUNT	68	0.007	0.077
CIMAREX	59	0.042	0.048
ANADARKO	53	0.044	0.040
DEVON	31	0.061	0.006
BP PRODUCTIONS	30	0.000	0.035
RANGE PRODUCTION	30	0.047	0.012
ALL OTHERS	669	0.446	0.565

within a firm should reveal no difference between auction and negotiations. Table 9 reports estimates of our main regressions, with and without fixed effects for the identity of the firm that wins the lease. Even after conditioning on firm identity, bonus payments, output and lease revenue are all still larger, by a similar magnitude, on auctioned leases than negotiated leases. If anything, these within-firm comparisons are even larger than than our baseline comparisons.

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

	Bonus	Bonus	Output	Output	Revenue	Revenue
Auction	0.36 (0.08)	0.39 (0.11)	0.107 (0.049)	0.182 (0.106)	3.52 (2.05)	5.30 (3.41)
Firm FE	No	Yes	No	Yes	No	Yes
N	1,274	1,274	1,175	1,175	1,175	1,175
R^2	0.953	0.970	0.752	0.831	0.766	0.854

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 controls for royalty rate, term, and a spline in acres.

Given that the differences between auctions and negotiations exist in comparisons within the same firm, we conclude that the source of the output effect must be due to horizontal

differences, or “match.” How plausible are idiosyncratic match shocks, which vary across potential firm-lease combinations, as a determinant of differences between auctioned and negotiated leases? While we are not aware of a direct test for this hypothesis, we can use the auction bid data to verify that the magnitude of firm-lease shocks must be large, relative to vertical differences among firms. If a firm’s value for a parcel was mostly vertical, in the sense that some firms were inherently more productive than others, we’d expect to see a consistent ranking of bids between firms, across auctions. In particular, when two firms bid in the same set of auctions, we’d expect the higher productivity firm to bid more than the lower productivity firm in every auction. 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 higher than the later firm (Firm B). If vertical differences across firms were more important than firm-lease match, we’d expect to see that one firm consistently bids higher than the other. What we observe is the exact opposite: for 11 of the 13 pairs, the fraction of the time that one firm wins more than the other is statistically identical to a coin toss.

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	CONOCO PHILLIPS	19	0.79	0.019
CIMAREX	RESOLUTE	19	0.53	1.000
CONOCO PHILLIPS	ENERGEN	19	0.37	0.359
ENERGEN	RESOLUTE	19	0.42	0.648
COG	RANGE PRODUCTION	17	0.41	0.629
CONOCO PHILLIPS	RESOLUTE	17	0.53	1.000
CIMAREX	MARSHFIELD OIL AND GAS	12	0.67	0.388
ENERGEN	MARSHFIELD OIL AND GAS	12	0.67	0.388
299 PRODUCTION	THREE RIVERS	10	0.70	0.344
ADESCAPE	ENERGEN	10	0.00	0.002
ADESCAPE	SEMEION	10	0.30	0.344
ENERGEN	SEMEION	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.

7 Discussion

Mineral leases that are allocated by auction generate more seller revenue and are matched to more productive firms than otherwise identical leases that transact informally under the

Relinquishment Act. What features of the two transaction processes are responsible for these differences? Unfortunately, we cannot answer this question directly. While the auction process is comprehensively documented by an administrative body, with public records of all submitted bids on all potential transactions, 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.

One possibility is that transactions on RAL parcels 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 \(2017\)](#). In our setting, increased demand for GLO auctions would arise from the fact that they are centralized. State auctions are widely publicized and routinely held, whereas a central challenge for firms in acquiring negotiated acreage (both in RAL and private land writ large) is identifying which land is leasable, and performing title research to determine who actually owns it. It is thus likely that the latter mechanisms would result in fewer participants. Note that while reduced competition in a hypothetical RAL auction would generate a reduction in seller revenues by itself, the fact that match quality, as defined in [Section 6](#), also declines, suggests that the subset of bidders that participate in negotiations must exclude the highest value buyer with positive probability.

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 auctions generate in our setting suggests that 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, and, anecdotally, we know that many transactions conclude before other parties even have the opportunity to participate. 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. Relatedly, it seems intuitive that landowners would have a difficult time committing to (and executing) a more formal process. In the most extensive survey of private mineral rights owners to date, only 21% of lessors in Pennsylvania reported ever consulting with a lawyer before transacting.²⁷ Conversely, GLO rules require that all parcels to be auctioned be announced via public notice, with clearly posted reserve prices. The requirement that the lease go to the high bidder is codified in state law and easily enforceable and observable.

How feasible would it be for landowners to hold an auction? While it is possible that 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 Texas GLO now uses one such platform, EnergyNet.com, that explicitly advertises its availability to private landowners. Given our main treatment effect estimate in Table 3 implies a greater than 40% increase in bonus payments, this gain from using an auction appears to far exceed the cost.²⁸ In this specific context, it's also possible to imagine the Texas GLO performing these auctions on the surface owner's behalf, and presumably internalizing some scale economies while doing so.²⁹

7.1 External Validity

How generalizable are these results to the broader population of mineral leases on private land in the United States, which are also allocated in an informal, decentralized fashion? One possible concern about predicting that the returns to auctions would be similar in other locations is that the negotiations in our sample are particularly inefficient or uncompetitive. If that were the case, the true causal effects of auctions, relative to negotiations, in other

²⁷Survey 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].

²⁸Note that RAL landowners only have a 50% claim to the gain from auctions. So unless the state bore half the costs, the effective fee from the RAL landowners perspective would be 20%, which is still far below the estimated auction gain.

²⁹Indeed, 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.

mineral leasing settings would be smaller than the effects we estimate here.

We begin by noting that the auctions against which these negotiated leases are compared are not particularly competitive. In Appendix Table A.11, we tabulate the number of auctions with 1, 2, 3, 4, 5, or 6+ bidders, and within those groups, compute the average bonus payment per acre and the median reserve margin. More than half of all GLO auctions receive only 1 successful bidder, and this fact seems to be known to potential bidders, as auctions that do receive more bids have substantially higher winning bids. The fact that reserve margins are much lower for the vast majority of auctions with 1 or 2 realized bidders, relative to auctions with more, suggests that either GLO has set reserve prices relatively low or that bidders expect a low, but positive probability of competition, a phenomenon studied in Kong (2017).

Similarly, it is unlikely that RAL negotiations are especially bad. Although data on the quality of negotiations in other settings is hard to come by, what little information is available suggests that private landowners are not particularly savvy. For example, the aforementioned Pennsylvania survey found that 79% of lessors only spoke to one E&P company before signing a lease. They also appear relatively uninformed, with only 32% reporting to have consulted any educational materials prior to signing.

In contrast, 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 shorter primary term. 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, so the treatment effects we estimate are likely to be lower bounds on the difference in revenues and allocative efficiency we would expect from replacing informal negotiations with centralized auctions in the broader private leasing population.

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 36 more log points in up front payments, and that auctions produce 46 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 (the Texas GLO), 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*, Forthcoming.
- Athey, Susan, Julie Tibshirani, Stefan Wager et al.**, “Generalized random forests,” *The Annals of Statistics*, 2019, 47 (2), 1148–1178.
- Backus, Matthew, Thomas Blake, Bradley Larsen, and Steven Tadelis**, “Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions,” Working Paper 24306, National Bureau of Economic Research February 2018.
- , **Tom Blake, and Steven Tadelis**, “Cheap Talk, Round Numbers, and the Economics of Negotiation,” June 2015, (21285).

- 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,” Working Paper 23060, National Bureau of Economic Research January 2017.
- Bhattacharya, Vivek, Andrew Ordin, and James W Roberts**, “Bidding and Drilling Under Uncertainty: An Empirical Analysis of Contingent Payment Auctions,” 2018. working paper.
- 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**, “To Trade or Not to Trade: Oil Leases, Information Asymmetry, and Coase,” 2019. working paper.
- 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,” 2018.
- 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.
- 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.
- 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,” 2018. 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**, “Sequential Auctions with Synergy and Affiliation Across Auctions,” 2016. working paper.
- , “Selective Entry in Auctions: Estimation and Evidence,” Working Paper 2017.
- Larsen, Bradley**, “The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions,” Working Paper 20431, National Bureau of Economic Research August 2018.
- Li, Fan, Kari Lock Morgan, and Alan M Zaslavsky**, “Balancing covariates via propensity score weighting,” *Journal of the American Statistical Association*, 2018, 113 (521), 390–400.
- 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.
- 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,” SSRN Scholarly Paper ID 2961795, Social Science Research Network, Rochester, NY May 2017.
- 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.

Timmins, Christopher and Ashley Vissing, “Environmental Justice and Coasian Bargaining: The role of race and income in lease negotiations for shale gas,” Technical Report, Working Paper 2017.


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

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 Parcel Characteristics Within the Leaseable RAL And Auction Sample

Table A.1 estimates similar specifications as Table 2, but is limited to RAL and State auction parcels that were leaseable as of 2005. This is the set of parcels underlying our sample leases. It also includes a variety of other surface characteristics that we use as “extra” covariates in the lease regressions.

Table A.1: Parcel comparison: Leasable Auction and RAL Land as of January 1, 2005

	Thickness	Acres	Shape	Water	Rivers	Roads	Shrub	Forest	Cultivated	Developed
Auction	-0.029 (0.061)	-0.069 (0.019)	-0.007 (0.006)	0.072 (0.386)	0.042 (0.089)	-0.038 (0.155)	-1.810 (1.050)	0.201 (0.739)	0.134 (0.575)	0.238 (0.155)
Average	3.488	0.290	0.960	12.084	0.587	1.866	77.359	6.025	2.052	0.733
N	1,202	1,763	1,763	1,763	1,763	1,763	1,763	1,763	1,763	1,763
R^2	0.895	0.454	0.469	0.824	0.531	0.426	0.913	0.793	0.627	0.897

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 in thousands, 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, Rivers is the distance in thousands meters to the nearest river or stream, Roads is the distance in thousands of meters to the nearest road, and developed high and low, cultivated and forests are land cover measures listed as fractions. 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.

A.3 Additional Bonus and Output Regression Results

Table A.2: Bonus Payments and Mechanism Type, per Acre

	(1)	(2)	(3)	(4)	(5)
Auction	0.78 (0.20)	0.83 (0.28)	1.02 (0.47)	0.75 (0.25)	0.92 (0.13)
Grid	10	10	10	20	DML
Time	Q	GY,Q	GYQ	GY,Q	DML
N	1,274	1,274	1,274	1,274	1,274
R^2	0.711	0.869	0.891	0.789	

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. \(2018\)](#). All models include a spline in acres and linear terms for term length and royalty rate.

Table A.3: Bonus Payments and Mechanism Type, per Acre: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.83 (0.28)	0.91 (0.12)	0.84 (0.28)	0.87 (0.13)	0.98 (0.49)	1.16 (0.19)
Estimate	G10Y	DML	G10Y	DML	G10Y	DML
Surface Controls	Yes	Yes	Yes	Yes	No	No
Thickness Controls	No	No	Yes	Yes	No	No
Private Only	No	No	No	No	Yes	Yes
N	1,274	1,274	1,070	1,070	1,073	1,073
R^2	0.870		0.863		0.880	

The dependent variable in each regression is the lease’s bonus payment per acre. Surface controls include shape regularity, a dummy variable for whether the lease spans multiple parcels, surface cover measures, and distance to roads and water sources. In columns 3 and 4, the sample excludes leases overlying parts of the Eagle Ford shale and all leases overlying the Haynesville and Barnett shales, for which there is no thickness information available. In columns 5 and 6, the sample is restricted to leases with private surface ownership. Columns 1, 3, and 5 use fixed effects for year-by-10-mile grid, as well as quarter of sample, with standard errors clustered by grid. Columns 2, 4 and 6 use a double/debiased Machine Learning routine to control for location and time, as recommended in [Chernozhukov et al. \(2018\)](#). All models include a spline in acres, and linear terms in term length and royalty rate.

Table A.4: Log Total Seller Revenue and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.35 (0.10)	0.36 (0.13)	0.44 (0.24)	0.33 (0.13)	0.51 (0.09)	0.36 (0.13)
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	1,175	1,175	1,175	1,175	1,175	974
R^2	0.687	0.856	0.913	0.753		0.853

The dependent variable is the natural logarithm of the discounted present value of the sum of bonus payments, delay rentals paid, and production royalties. 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. Column 5 uses a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). All models include a spline in acres, and linear terms in term and royalty rate. “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 the thickness of the shale formation. The sample includes all leases whose primary term ends before March, 2019.

Table A.5: Log Bonus Results on the Poisson Samples

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.33 (0.06)	0.35 (0.10)	0.41 (0.16)	0.37 (0.08)	0.49 (0.05)	0.38 (0.09)
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	1,093	738	613	944	1,175	618
R^2	0.855	0.951	0.971	0.910		0.953

The dependent variable is the natural logarithm of bonus payments, per acre. In columns 1-4 and 6, 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. Column 5 uses a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). For estimation details, see [A.7](#). All models include a spline in acres, and linear terms in term length and royalty rate. “Extra” controls include shape regularity, a dummy variable for whether the lease spans multiple parcels, surface cover measures, and distance to roads and water sources. In each column, the sample is identical to the corresponding column in the bottom panel of [Table 6](#).

Table A.6: Log Total Seller Revenues Results on the Poisson Output Sample

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.33 (0.10)	0.39 (0.18)	0.60 (0.33)	0.32 (0.13)	0.51 (0.09)	0.38 (0.18)
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	1,093	738	613	944	1,175	618
R^2	0.662	0.767	0.824	0.697		0.768

The dependent variable is the natural logarithm of the discounted present value of the sum of bonus payments, delay rentals paid, and production royalties, per acre. In columns 1-4 and 6, 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. Column 5 uses a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). For estimation details, see [A.7](#). All models include a spline in acres, and linear terms in term length and royalty rate. “Extra” controls include shape regularity, a dummy variable for whether the lease spans multiple parcels, surface cover measures, and distance to roads and water sources. In each column, the sample is identical to the corresponding column in the bottom panel of [Table 6](#).

Table A.7: Drilling

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.062 (0.040)	0.056 (0.047)	0.098 (0.075)	0.070 (0.054)	0.077 (0.038)	0.075 (0.048)
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	1,175	1,175	1,175	1,175	1,175	974
R^2	0.407	0.672	0.761	0.518		0.660

The dependent variable is equal to 1 if a lease is drilled and 0 otherwise). In columns 1-4 and 6, 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. Column 5 uses a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). Standard errors are clustered by grid in columns 1-4, and 6. All models include a spline in acres, and linear terms in term length and royalty rate. “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 the thickness of the shale formation. The sample includes all leases whose primary term ends before March, 2019.

A.4 Overlap-weighted ATEs

In Section 3, we describe our rationale for dropping (primarily) negotiated leases with especially large sizes or terms that are shorter than 1 year. Our goal in dropping these observations is to achieve balance between the observable characteristics of auctioned and negotiated leases. Here we report the results that we obtain from measuring the causal effects of mechanism type on lease outcomes using overlap-weighted treatment effect techniques described in [Li et al. \(2018\)](#). These estimates do not require *a priori* choices about which comparisons are sufficiently balanced, so the data for these calculations include leases we previously dropped because they were too large or their terms were too short. We compute the estimates in Table A.8 using the “causal forest” estimator from the `grf` R package, as described in [Athey et al. \(2019\)](#). For each outcome, these results have the same sign as those in in Tables 3 and 6, though the magnitudes are larger.

Table A.8: Overlap-weighted ATE Estimates

	Bonus	Bonus	Output	Output	Revenue	Revenue	Drilled	Drilled
Auction	0.56 (0.05)	0.52 (0.05)	0.236 (0.064)	0.232 (0.070)	13.15 (4.38)	14.05 (5.54)	0.100 (0.037)	0.068 (0.035)
Extra	No	Yes	No	Yes	No	Yes	No	Yes
N	1,595	1,282	1,467	1,164	1,467	1,164	1,467	1,164

Overlap weighted average treatment effect estimates for the natural logarithm of the bonus payment (columns 1 and 2), Output per acre (columns 3 and 4), Revenue per acre (columns 5 and 6) and whether the lease is drilled (columns 7 and 8). Columns 3 through 8 are estimated using the subsample of leases whose primary terms are not censored. The covariates used in determining the conditional expectation and propensity functions include location, effective date, acres, term, royalty rate, as well as shape regularity, surface cover, an indicator for leases spanning multiple parcels, shale thickness and distance to roads and water sources for the “extra” columns 2, 4, 6, and 8. Shale thickness is not available for leases in parts of the Eagle Ford shale nor for any leases in the Haynesville and Barnett shales.

A.5 Delay Rental Statistics

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. Failure to drill a productive well before such milestones or to make a required delay rental payment result in early termination of the lease. For example, a lease may have a 5 year primary term, a 25% royalty rate, a bonus payment of \$1,000 per acre, and delay rentals of \$200 per acre due at the start of the fourth year. Here, we estimate the causal effect of transaction mechanism type on the magnitude of potential delay rentals in a lease, as well as the portion of those contracted rentals that are ultimately paid.

Table A.9: Potential Delay Rentals and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	-0.07 (0.06)	-0.10 (0.10)	0.00 (0.17)	-0.09 (0.05)	-0.09 (0.03)	-0.11 (0.10)
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	1,274	1,274	1,274	1,274	1,274	1,070
R^2	0.556	0.716	0.819	0.648		0.721

The dependent variable in the first two columns is potential discounted delay rentals, in thousands of dollars per acre. The specifications in columns 1, 3, and 5 control for location and time using 10-mile grid by year-quarter of sample fixed effects, with standard errors clustered by grid, while columns 2, 4 and 6 use a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). All models include a spline in acres and linear terms for term length and royalty rate.

Table A.10: Delay Rentals Paid and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	-0.07 (0.03)	-0.08 (0.06)	-0.04 (0.09)	-0.08 (0.05)	-0.07 (0.02)	-0.08 (0.06)
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	1,175	1,175	1,175	1,175	1,175	974
R^2	0.371	0.593	0.754	0.515		0.598

The dependent variable in the first two columns is discounted delay rentals in thousands of dollars per acre, that were actually paid. The specifications in columns 1, 3, and 5 control for location and time using 10-mile grid by year-quarter of sample fixed effects, with standard errors clustered by grid, while columns 2, 4 and 6 use a double/debiased machine learning routine, as recommended in [Chernozhukov et al. \(2018\)](#). The sample includes all leases whose primary term ends before March, 2019. All models include a spline in acres and linear terms for term length and royalty rate.

A.6 Auction Statistics

Table A.11: Auction Results by Number of Bidders

Bids	Auctions	Fraction	Avg. Bonus (\$/Acre)	Med. Bid/Reserve	Med. Markup
1	235	0.55	1,337	1.62	
2	88	0.21	1,945	2.08	1.39
3	35	0.08	3,014	3.61	1.40
4	30	0.07	2,743	3.26	1.12
5	21	0.05	6,418	3.61	1.10
6 +	19	0.04	6,027	4.71	1.14

This table summarizes the winning bids from GLO auctions. The fifth column contains the median ratio of winning bid to the reserve price for each group, while the last column shows the median ratio of winning bid to second highest bid.

A.7 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-validation, and report the “median” values of these across 101 randomized cross-validation 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. 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. Note that this objective only allows for a single binary covariate in D , so in the DML pseudo-poisson regression in Table 6, we include lease size, royalty rate and term length in X .

Appendix B Data Cleaning

B.1 Sample construction

Table A.12 presents the number of RAL and auction leases from original data provided to us by GLO that survive each of our sample restrictions. We use the EIA’s definition of shale formations in Texas, shown shaded in yellow in Figure A.1. “Undivided” ownership refers to shared ownership of real property by two or more parties.³⁰ A small number of non-RAL leases are actually allocated via bilateral negotiation when it is situated in a position where only one party can economically use it. Similarly, some RAL leases are allocated by auction when the State is unable to determine who the rightful surface owner is.

Table A.12: Sample Construction

	Drop Reason	Negotiation	Auction
All Leases		3,977	740
	Not on Shale	2,350	528
	Missing Value	2,305	528
	Less Than 10 or Greater Than 1,000 Acres	1,832	476
	Gross and Net Acreage Differ	1,327	474
	Undivided Interest	870	460
	Term Less Than 1 Year	860	460
	Cancelled or Withdrawn	852	457
	Negotiated State Lease	847	457
	Auctioned RAL Lease	847	428
	Lessee Owns RAL Surface	846	428
Final Sample		846	428

Additional discussion provided in Section 3.1.

³⁰For 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.

B.2 Firm Names

Though we observe the name of the firm on the lease, E&P companies sometimes use intermediaries 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 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 RAL Lease Addenda

In addition to specifying a bonus payments, royalty rate 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.

To determine the validity of this concern, 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. About 73% of RAL leases have one or more additional clauses in their lease addenda. In Table A.13, we include measures of these protective clauses in bonus regressions like those shown in Table 3. The first two columns mirror the result shown in the main text: auctioned leases pay 36 to 58 more log points in up-front bonus payments than negotiated leases do. 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 bonus payments.

Table A.13: Bonus Payments and Mechanism Type: Robustness to RAL Lease Addenda

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.36 (0.08)	0.52 (0.05)	0.40 (0.09)	0.58 (0.05)	0.39 (0.09)	0.57 (0.05)
Pages			0.00 (0.01)	0.01 (0.01)		
Clauses			0.01 (0.00)	0.01 (0.00)		
Surface Protection					0.03 (0.04)	-0.05 (0.05)
Payment Terms					0.01 (0.04)	0.04 (0.03)
Location Requirements					-0.05 (0.07)	-0.02 (0.07)
Pugh Clause					0.08 (0.04)	0.10 (0.05)
Cleanup Terms					0.05 (0.05)	0.14 (0.08)
Livestock Protection					-0.19 (0.12)	-0.16 (0.09)
On-site Water Use					-0.09 (0.05)	-0.10 (0.05)
Waste Management					0.03 (0.07)	-0.19 (0.10)
Definitional Changes					0.02 (0.03)	0.01 (0.04)
Pollution Protection					-0.08 (0.11)	0.15 (0.10)
Infrastructure Constraints					0.12 (0.06)	0.04 (0.06)
Caliche Use					-0.14 (0.12)	0.20 (0.11)
Additional Fees					-0.03 (0.05)	0.09 (0.07)
Time Constraints					-0.01 (0.15)	-0.11 (0.10)
Miscellaneous					0.05 (0.06)	0.20 (0.07)
Grid	10	DML	10	DML	10	DML
Time	GY,Q	DML	GY,Q	DML	GY,Q	DML
N	1,274	1,274	1,274	1,274	1,274	1,274
R^2	0.953		0.953		0.954	

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 and effective date.