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

We thank participants at numerous seminars, as well as Peter Maniloff, Aaron Flaaen, Timothy Fitzgerald and Piotr Dworczak 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.

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
JEL No. D44,L13,Q35

ABSTRACT

This paper compares outcomes from informally negotiated oil and gas leases to those awarded via centralized auction. We use data on all contractual characteristics and production outcomes for a class of state-owned mineral rights overlying newly discovered shale formations in Texas, between 2005 and 2016. On roughly three quarters of this land, the Texas Relinquishment Act of 1919 authorizes private individuals who own surface-only rights to negotiate mineral leases on behalf of the public in exchange for half of the proceeds. The remainder are allocated via centralized auctions. Using variation from this natural experiment, we find that almost a century after leasing mechanisms were assigned, auctioned leases generate 67% larger up-front payments than negotiated leases do. The two mechanisms also allocate mineral rights to different oil and gas companies, and leases allocated by auction are 44% more productive. These results are consistent with theoretical intuitions that centralized, formal mechanisms, like auctions, outperform decentralized and informal mechanisms, in both seller revenues and allocative efficiency. Our findings have important implications for the more than $3 trillion of minerals owned by private individuals in the US, the vast majority of which transact in informal and decentralized settings.

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 possible partners. Many land transactions also look like this. How should an owner go about this process? The fact that buyer valuations and even the identities of interested buyers may not be known to sellers provides the starting point for rich literatures investigating theoretically optimal mechanisms or search processes. However, there is little evidence on the consequences of less structured, and likely suboptimal, allocation mechanisms used in the real world.

In this paper, we directly measure the gains from using a centralized, theoretically high-performing mechanism, relative to using an informal, decentralized mechanism, in the market for mineral leases in Texas. We focus on a large class of lands 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] Although the minerals within these lands belong to the State, on some parcels, the Texas Relinquishment Act of 1919 grants the private individuals who separately own the surface 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 similar, 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 more than thirteen hundred auctioned and negotiated leases for publicly owned minerals in Texas between 2005 and 2016, we find that auctioned leases sell for $709 per acre more than negotiated leases do. This precisely estimated difference is 67% of the average price of a negotiated lease, so the causal effect of auctions on seller revenues is economically important. Furthermore, this result is robust to a wide range of controls and sample restrictions, is not caused by differences in the likelihood that auctioned and negotiated parcels have a successful transaction, and even persists after

[1] Throughout 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.

conditioning on the resulting *ex post* investment and output. For the average negotiated lease in our data, the difference in seller revenues is worth about $202,000.

In addition to this improvement in seller revenues, which in principle could simply reflect a transfer from buyers to sellers, we also look for evidence of efficiency gains. 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 measures of investment and output across auctioned and negotiated leases. In our data, auctioned leases are 22% more likely to be drilled and produce 44% 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 $249,000 per lease. Moreover, we show that while auctions allocate minerals to different firms, the payment, investment, and output results all hold *within* firm, suggesting an important role for firm-lease "match" in determining output. As we discuss in section 7, the importance of match-based differences in productivity across potential lessees suggests limited scope for "nudge" style policy interventions which do not alter the mechanism used to improve outcomes in decentralized markets.

Our comparison of auctions to the unobserved distribution of informal mechanisms provides an empirical analogue to a rich theory literature comparing one formal mechanism to another (Milgrom, 2004; Klemperer, 2004). Much of this work compares an auction to some multi-stage mechanism either designed to highlight a feature or flaw of auctions or inspired by the norms of a particular marketplace. Bulow and Klemperer (1996) shows that sellers gain more by adding an additional bidder to an otherwise unoptimized auction than they do by having all the bargaining power, perfect information and an optimally designed two-stage mechanism. One interpretation of this result, consistent with our empirical findings, is even if the unobserved mechanisms embedded in our "negotiations" are as seller-optimal as they can possibly be, centralized auctions will still perform better if they are able to attract more competition.

Subsequent work sought to incorporate entry costs into this comparison, and considered a sequential mechanism inspired by real-world corporate takeover battles. Bulow and Klemperer (2009) show that while sequential transaction mechanisms can generate total welfare gains, relative to auctions, by avoiding the excessive entry in auctions, these gains mainly accrue to buyers, to the point that sellers should still prefer auctions. Roberts and Sweeting (2013) show that when entry is selective, so that high value bidders are more likely to participate, sequential mechanisms can outperform auctions, even those with optimally set reserve prices. They empirically verify these results using a structural model of auctions with endogenous entry in the market for timber. Wang (2018) arrives at a similar result, in a model of seller-optimal mechanisms for corporate takeovers. Inspired by the recent growth of "go-shop" sequential mechanisms in these markets, she shows that, depending on parameter values, the seller-optimal mechanism may involve sequential negotiations, a simultaneous auction, or a take-it-or-leave-it offer to a specific potential buyer.

As we discuss in section 7, the fact that we do not observe the details of the "mechanisms" employed in negotiations means that our results cannot empirically resolve any (surely context specific) theoretical ambiguity on the "optimal" way to sell something. Conversations with industry participants suggest that features of the proposed sequential mechanisms from the theory literature coexist with costly landowner search effort (Hortaçsu and Syverson, 2004; Allen et al., Forthcoming; Cuesta and Sepulveda, 2018), bilateral bargaining (Backus et al., 2015, 2018; Larsen, 2014) and even some take-it-or-leave-it behavior on the part of some buyers. Rather, the contribution of this paper is to demonstrate the magnitude of the gains from using a fairly standard, nearly optimal mechanism, in a real world setting.

As such, this paper joins a small, but growing, empirical literature that compares the performance of real-world, non-auction mechanisms, to auctions.[2] Larsen (2014) shows that bilaterally negotiated used car transactions that follow failed auctions generate a large fraction of the theoretical maximum available *ex ante* surplus, suggesting that some commonly used real-world mechanisms have performance comparable to optimal auctions. Salz (2017) documents large inefficiencies in the highly decentralized market for waste collection in New York City, and finds that intermediaries who perform procurement auctions on behalf of their clients have spillover effects which reduce the costs of all buyers, including those who do not use intermediaries. Finally, Gentry and Stroup (2018) estimates a model of corporate takeover auctions. In each of these papers, only one mechanism is observed in the data. To predict what would happen in 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 Larsen (2014), its the optimal mechanism in Myerson (1981); in Salz (2017) its a pure search market; in Gentry and Stroup (2018) its a sequential mechanism with possibly selective entry). In contrast, we observe the results of auction and non-auction mechanisms simultaneously, on otherwise identical objects. As a result, we can directly compare welfare relevant outcomes across mechanisms.

We also contribute to the large literature on the economics of oil and gas leasing in the United States. Early work by Ken Hendricks and Rob Porter on the performance of auctions

[2] There 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 an 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.

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). In symmetric information environments (auctions for "wildcat" tracts), they showed that US government auctions captured approximately 100% of the *ex ante* surplus, while it captured considerably less in asymmetric information environments (auctions for "drainage" tracts). In contrast to this literature on the performance of sellers in the presence of informed vs. uninformed bidders, in this paper we analyze seller performance under formal vs. informal mechanisms. Recent work on mineral lease auctions has sought to separately identify affiliation from synergies between neighboring parcels, and to measure the extent to which bidders have uncertainty about how many competitors they will face (Kong, 2017, 2016). Bhattacharya et al. (2018) examine the economic structure of the "security" sold by the winning bidder to the auctioneer, and find that the New Mexico State Land Office's choice of this security is close to optimal, relative to their estimates of the structural demand and cost primitives. To our knowledge, our paper is the first to document the revenue and allocative efficiency properties of non-auction mechanisms in mineral leasing.

In addition to quantifying differences in investment and output between negotiated and auctioned leases, we also document for the first time the *exact level* of revenues that private landowners earn from oil and gas development. Approximately three quarters of all mineral rights in the United States are held by private individuals. However, the vast majority of leases on these lands are incompletely recorded in county registries. In particular, it is common for E&P companies to record a placeholder value for bonus payments in their public lease filings, often *"$10 plus other consideration paid."* Moreover, most mineral leases are developed as a part of larger "units" in which several leases are combined into a single project. This "unitization" creates a challenge for matching publicly documented drilling and production data to individual mineral owners. As a result, little is know about total landowner revenues in mineral rights contracts. In our setting, we observe the full set of payments received by the State and private surface owners, including bonus payments and royalty revenue, even on unitized leases. Previous literature on the landowner benefits of the fracking boom was not able to directly measure revenues in this way (Brown et al., 2016; Feyrer et al., 2017; Bartik et al., 2017). We find that bonus payments represent 80% 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. Precisely quantifying these landowner benefits provides important context to the growing number of papers documenting large local negative externalities from fracking (Muehlenbachs et al., 2015; Currie et al., 2017).

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

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.[3] 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.[4] 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. However, 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. Many leases also include clauses regarding operations, cleanup and other landowner protections (Vissing, 2017).[5]

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."[6] One reason that E&P companies use landmen is that

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

[4]This 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).

[5]We study these "lease addenda" formally in Appendix C.

[6]Landmen are also used in surface rights transactions, such as siting pipelines or wind farms.

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.

2.2 Texas Relinquishment Act

Private mineral rights are a uniquely American phenomenon. When individuals outside of the US 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. Eventually, 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 subsequently sold to private citizens, Texas, following in Spanish 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 mere $0.10 per acre annual fee.[7]

When oil was discovered in Texas at the turn of the 20th century, many surface owners of PSF land argued that this compensation was inadequate.[8] To stave off "armed rebellion" by the surface owners against state lessees, the Texas legislature passed the Relinquishment Act of 1919 (Shields, 1981). This law, amended and reinterpreted through a decade and a half of subsequent litigation, appointed the surface owner as the minerals leasing agent of the state, provided that the surface owner's parcel had been acquired from the PSF by 1931.[9] In exchange for negotiating a lease on the state's behalf, surface owners were awarded half of bonus and royalty payments generated from their land.

Following the passage and eventual legal settlement of the Relinquishment Act in 1931, the mineral status of land sold from the PSF took on two forms. Transactions between 1931 and 1973 explicitly awarded the majority of minerals to new surface owners. On

[7]Interestingly, when Texas declared independence, it *released* full ownership of minerals on parcels privatized during colonial control to surface owners.

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

[9]The Texas Supreme court finalized their interpretation of the Relinquishment Act in mid 1931.

these "Free Royalty" lands, the state retained just a $1/16^{\text{th}}$ royalty interest, and no bonus interest. Starting in 1973, the state explicitly retained all mineral interests in subsequent land transactions. As a result, land in the initial PSF belongs in one of four categories: land governed by the Relinquishment Act (which we refer to as "RAL" parcels), with a 50/50 split of the mineral interest between the surface owner and the State; land sold between 1931 and 1973, in which the State retains a minority royalty interest; land sold after 1973, in which the state owns the entirety of the mineral interest; and land not yet sold, in which the state owns the entirety of the surface and minerals. Our primary interest in this paper is a comparison of leasing activity on Relinquishment Act lands with leasing activity on other PSF land not sold as of 1973. Leases on these types of land are managed by the Texas General Land Office, with rigorous and publicly observable record keeping of all contractual terms. Since leases on Free Royalty Land are not managed by the GLO, bonus payments on theses leases are not publicly recorded.

When a surface owner of an RAL parcel is approached by an E&P company, she initially negotiates a lease in the same fashion as she would if she owned a parcel with private mineral rights. However, once the surface owner and the E&P company reach an agreement, they must submit their lease to the GLO for final approval. If approved, the lessee remits half of the bonus and royalty payments to the state. For PSF land sold after 1973 or not yet sold, the state directly leases mineral rights to E&P companies. However, unlike leases on RAL parcels, or the broader population of private leases, the state awards leases on these non-RAL parcels using an auction. In these auctions, 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. The state awards the lease to the highest bidder at the bonus payment that bidder submitted.[10] In the event that an E&P company wishes to sign a lease in an RAL parcel, and the surface owner cannot be found, the GLO will conduct an auction for the lease using its standard process and the surface owner's share of the proceeds either go into escrow or are forfeited to the State of Texas. As a result, there are a small number of auctioned leases on RAL parcels, which we include in our sample of auctioned leases below.

[10]Similar processes are used in state land offices outside of Texas and at the Federal level.

3 Data

Our primary dataset consists of the universe of oil and gas leases signed on PSF land in Texas between 2005 and 2016.[11,12] Our initial dataset includes the shape, location, size, effective date, bonus payment, primary term and royalty rate for 4,076 RAL leases, 111 of which are allocated by auction, and 851 State leases. For all leases that eventually result in drilling, we observe monthly payments for gas and oil royalties remitted to the state, up through June, 2018. We combine this with royalty rate and output price[13] information to infer which leases were drilled and 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 which performed the title research on Relinquishment Act lands for the State of Texas. P2Energy Solutions determined the time at which each parcel in this map was initially privatized, or if it has not yet been privatized. We use this map of parcels to characterize differences in the likelihood of a successful lease across negotiation (RAL) and auction (non-RAL) parcels.

GLO uses first price, sealed bid auctions to allocate its non-RAL leases.[14] For every parcel that is nominated by an E&P company for inclusion in the next GLO auction, we observe a "bid notice" describing the parcel itself, the date that the auction will be held[15], and the reserve price. Following the auction, we observe the name of each bidder who bid above the reserve as well as their bid. We infer that a parcel received no bids above the reserve price if it exists in a bid notice but does not subsequently show up in the post-auction bid data.

[11] In addition to the PSF, the GLO also manages mineral leasing for land owned by several "Land Boards", including those for Veterans, Parks & Wildlife, the Department of Criminal Justice, as well as land underneath state roads, waterways and water bodies. Leases for this land are also allocated by auction and are included in the GLO's public lease data. However, we do not use it in this project, as the "assignment" to the auction treatment is not quasi-random, as in the PSF.

[12] We also collected data from the universe of parcels from another public entity in the Texas, the Texas Board for Lease of University Lands. However, while leases on University Lands are allocated in centralized auctions, they are not directly comparable leases on RAL parcels, because the parcels are much larger and more contiguous. University Lands parcels are also infrequently situated near RAL parcels. Finally, University Lands parcels have been acquired over time, and do not fit the clean quasi-experimental comparison with leases on RAL parcels discussed in Section 4.

[13] We use the Henry Hub price for natural gas royalties and the WTI price for oil royalties.

[14] Note that in Texas, these procedures are formally called "sealed bid sales" and are never referred to as auctions, in an effort to distinguish them from a prior era of open outcry English auctions.

[15] 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.

3.1 Data cleaning and sample selection

We impose a number of restrictions on these data to obtain our final sample.[16] 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.[17] Additionally, we exclude leases that have missing values for one or more contractual characteristics (bonus, royalty, primary term, size or location), leases with undivided[18] ownership issues, leases whose GIS shapes are inconsistent with their stated size, leases smaller than 10 acres or bigger than 1,000 acres[19], leases with primary terms shorter than 12 months, a small number of non-RAL leases that were actually allocated by negotiation[20], and a smaller number of leases acquired by the PSF after its initial allocation. The resulting dataset of 860 negotiated leases and 460 auctioned leases is summarized in Table 1. Figure 1 demonstrates the distribution of lease types over time.

Figure 1: Sample Leases by Year and Type

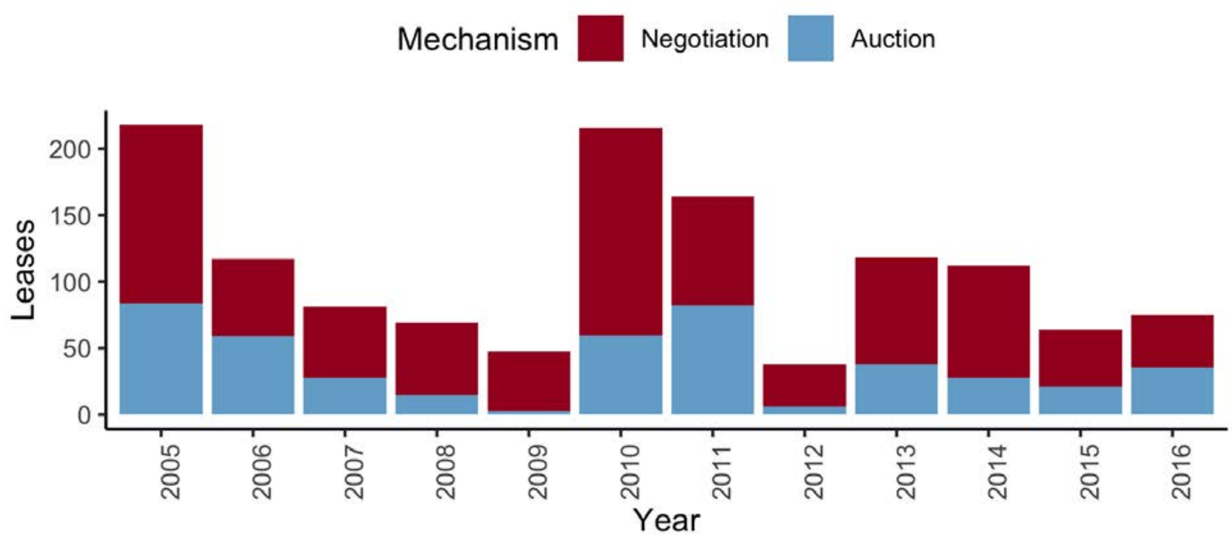

In the cross section, after imposing these sample restrictions, auctioned leases are larger, have slightly "more convex" shapes, and are less likely to cover more than one legally defined

[16]Table A.6 shows how many of the initial observations survive each of these cuts.

[17]We use the EIA's definition of shale formations in Texas, shown shaded in yellow in Figure A.1

[18]"Undivided" ownership refers to shared ownership of real property by two or more parties. For 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.

[19]Individual parcels we observe in the Texas Permanent School Fund are never more than 1,000 acres, and GLO rarely auctions leases that cover more than one parcel.

[20]Some non-RAL land is allocated via bilateral negotiation when it is situated in a position where only one party can economically use it.

Table 1: Lease Summary Statistics by Type

	Negotiation (N = 860)				Auction (N = 460)					
Variable	mean	sd	min	max	mean	sd	min	max	Difference	p-value
Land Characteristics										
Acres	0.29	0.26	0.01	1.00	0.35	0.25	0.01	0.77	-0.07	0.00
ShapeQuality	0.94	0.14	0.10	1.00	0.96	0.10	0.28	1.00	-0.01	0.06
MultiPolygon	0.05	0.22	0.00	1.00	0.03	0.16	0.00	1.00	0.02	0.02
Lease Characteristics										
Bonus	1.06	1.44	0.03	26.84	2.04	2.47	0.02	15.12	-0.98	0.00
Term	46.24	13.88	12.00	60.00	57.18	7.73	36.00	60.00	-10.94	0.00
RoyaltyRate	0.24	0.02	0.19	0.25	0.25	0.01	0.20	0.25	-0.01	0.00
Lease Outcomes										
Drilled	0.37	0.48	0.00	1.00	0.32	0.47	0.00	1.00	0.05	0.07
Output	0.18	0.44	0.00	4.01	0.18	0.39	0.00	2.37	0.00	1.00
Revenue	2.61	4.37	0.03	42.17	3.65	4.54	0.02	33.59	-1.04	0.00

Units: acres are reported in thousands; bonus and revenue 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 months. *Definitions*: shape quality is the ratio of the lease's size to the size of the convex hull containing it; "MultiPolygon" leases have claims to minerals in multiple adjacent or disjoint parcels; we define a lease as "drilled" if it ever reports a royalty payment.

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, while auctioned leases have longer primary terms. Auctions are slightly less likely to be drilled, produce equivalent amounts of output, and the difference in total revenues (bonus payments plus royalty income) is similar to 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 these data to measure how formal auctions effect lease outcomes, including seller revenues, investment, and production, relative to outcomes on negotiations. In the ideal experiment, we would have randomized mechanism type, formal 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 on the parcels that were placed in the PSF by the Texas Constitution of 1876, and subsequently revealed to overlie shale formations in the early 2000s.[21] Within this sample, mechanism assignment is determined not by randomization, but by the date on which a parcel was subsequently privatized. Our auction "treatment" group contains leases on parcels that were not privatized before 1973, and as such are allocated using a formal auction. Our "control" group of informally negotiated leases lie in parcels that were privatized before 1931, when the Relinquishment Act was finalized.[22] Thus, to interpret differences in outcomes between these two groups as representing the causal effects of auctions, relative to informal negotiations, we must assume that classification into these two groups is uncorrelated with unmeasured determinants of lease outcomes during the recent shale boom.

To assess this assumption, note that we are able to directly control for the two primary

[21] As described in Section 3.1, we exclude areas of the PSF allocation which were not affected by the recent shale boom. However, in this section, we refer this subsample loosely as the "PSF".

[22] As discussed in Section 2, parcels in the missing group from this target population, so-called Free Royalty parcels sold between 1931 and 1973, are also informally negotiated. However, because the State of Texas only retained a $1/16^{\text{th}}$ *royalty* interest in these parcels, the entire bonus accrues to surface owners, which the GLO does not observe. As a result, although leases on these parcels are negotiated, we can't compare their bonus payments with the auction sample.

determinants of lease outcomes: *where* leases are and *when* they transact. Leases on parcels with better mineral resources may transact at higher prices, attract more investment and produce more output. Similarly, leases that occur during periods of high output prices or increased technological progress may earn higher prices or generate better outcomes. To ensure that differences in location and/or time do not confound our comparisons between the two lease types, we directly control for space and time using a set of location and time fixed effects or by estimating a non-parametric control function for location and time. Both of these strategies ensure we are making comparisons between leases with similar mineral quality which transact at similar times. We also condition on detailed information about how the surface is used and how far the parcel is from other potentially valuable features like water and roads. Thus, our identification assumption is that whether a parcel was sold prior to 1931 or later than 1973 is independent of any *residual* determinants of lease outcomes during the shale boom, after controlling for time, space, and other surface qualities.

The main threat to the validity of this assumption is the possibility that the forces which determine *when* a PSF parcel was privatized (if at all) are correlated with these residual unobserved determinants of lease quality in the current shale environment. For example, if land buyers in the pre-1931 era had knowledge about which parcels in narrowly defined areas would be better or worse producers in the modern shale environment, they might rationally have acquired the "good" parcels, 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 parcels would be worse, on average. Both of these scenarios require widespread knowledge about the precise location of the best mineral resources, several decades before shale formations were discovered or the technology to access them was invented. We view this as unlikely. Moreover, its worth noting that that when many of these RAL parcels were initially privatized, the State of Texas explicitly intended to retain full rights on any subsequent mineral revenues. The 50/50 split that exists in RAL leases today was only a consequence of mineral discoveries that occurred in other parts of the state, after the PSF rules were established.

Though we can't directly test whether our assumption that controls for location, time, and surface characteristics leave only "random" variation in outcomes, we can check whether parcels governed by the Relinquishment Act appear similar on observable dimensions to parcels subject to auctions. Table 2 presents a series of balance tests where we regress potentially confounding observable characteristics onto Auction_i and location fixed effects. Auction and negotiation parcels appear very similar, with no measurable differences in shape quality, land cover, or distance to infrastructure. The one exception is parcel size: auction parcels are smaller, and the difference is precisely estimated, even after conditioning on geographic location. Because of this, we control for lease size using a cubic spline in the lease's acreage with five knots in all of our regression specifications.[23]

Table 2: Parcel comparison

	Acres	Shape	Road Dist.	Dev. High	Dev. Low	Cultivated	Forest
Auction	-69.342	-0.007	-37.646	0.002	0.010	0.002	0.002
	(19.180)	(0.006)	(154.534)	(0.002)	(0.006)	(0.006)	(0.007)
N	1,747	1,747	1,747	1,747	1,747	1,747	1,747
R^2	0.451	0.467	0.422	0.304	0.415	0.633	0.793

Definitions: acres in thousands, shape quality is the ratio of parcel size to the size of the convex hull containing the parcel, road distance is the distance in meters from the parcel's centroid to the nearest road, developed high and low, cultivated and forests are land cover measures listed in percentage points. 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.

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

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

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 is the a spline in the lease's size in acres, its primary term in months, its royalty rate and, in some specifications, the "quality" of the lease's shape[24], a dummy variable indicating whether the lease covers more than one parcel, its distance to infrastructure, and satellite measures of its landcover characteristics. δ is a set of location and time fixed effects, location-by-time fixed effects, or a non-parametric function of location and time. Since 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, we estimate these models using several fixed effect specifications, as well as specifications that non-parametrically control for location and time by combining the Robinson (1988) transformation with modern random forest techniques (Athey et al., 2019).[25] τ thus reflects the difference in outcomes for auctioned leases relative to negotiated leases, within leases that have similar size, are located nearby, transact at similar times and have comparable

[23]Note that while negotiation *parcels* are larger than auction parcels, the reverse is true of sample *leases*.

[24]We define Shape Quality as the ratio of the lease's area to the area of the convex hull containing the lease.

[25]Specifically, we estimate these models by separately projecting Y, Auction and X onto the lease's GPS coordinates $L(i)$ and effective date $T(i)$ using random forests, and then regress the residuals $Y_i - \mathbb{E}[Y \mid L(i), T(i)]$ onto the residuals $\text{Auction}_i - \mathbb{E}[\text{Auction} \mid L(i), T(i)]$ and $X_i - \mathbb{E}[X \mid L(i), T(i)]$. We report heteroskedasticity robust standard errors for the Auction coefficient in these specifications, consistent with the results in Robinson (1988).

characteristics.

5 Seller Revenue Results

We begin by investigating the impact of auctions on seller revenues, estimating several versions of Equation 1 with bonus per acre as the dependent variable.

Table 3 presents the results. In column 1, we include fixed effects for the year-quarter of the lease's effective date and for the 10 square mile grid containing the lease's centroid. The interpretation of this estimate is that auctioned leases receive $584 more per acre in bonus payments than similar negotiated leases, a difference which is precisely estimated. This difference is 55% of the average per-acre bonus payment for RAL leases in Table 1. In column 2, we interact the grid indicators with year of sample indicators, to account for the fact that different locations in Texas were developed at different times in our sample. With these interactive fixed effects, the estimated auction coefficient is larger, still precisely estimated, and is 67% of the average negotiated bonus payment. This model, which compares leases for minerals that are located at roughly the same space 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)	(6)
Auction	584.52 (248.02)	708.85 (342.20)	1008.70 (490.63)	658.10 (256.45)	717.21 (160.32)	718.15 (340.91)
Grid	10	10	10	20	RF	10
Time	Q	GY,Q	GYQ	GY,Q	RF	GY,Q
Extra	No	No	No	No	No	Yes
N	1,320	1,320	1,320	1,320	1,320	1,320
R^2	0.584	0.760	0.895	0.677	0.046	0.763

The dependent variable in each regression is the bonus payment ($) 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 Robinson (1988) transformation using random forests, with heteroskedasticity-consistent standard errors. All models include a spline in acres and linear terms for 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 average negotiated bonus payment is $1,060 per acre.

In the remaining columns we investigate the sensitivity of these results to the inclusion of additional controls[26]. In column 3, we include location-quarter-of-sample fixed effects to impose more stringent limits on which leases can be compared over time. This results in an even larger, and still-precise estimate: at $1,009 per acre, the difference is on par with the average negotiated bonus payment. To ensure that our results are robust to different choices of spatial controls, in column 4 we use 20 square mile grids instead of 10 square mile. 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. Finally, in column 6, we include controls for other measures of surface quality, like the quality 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. Across all of these specifications, we find consistent evidence that bonus payments are substantially larger in auctions than they are in negotiations.[27]

Even at the lower end of these estimates, the implications for seller revenue are large. For an RAL lease of average size, the causal effect of auctions on seller revenues is worth about $167,000. In aggregate, this increase in RAL bonus payments would be worth about $143 million in our sample.

5.1 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 bonus payments, royalties, primary terms, or additional protective clauses in leases 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. Thus, if failures are common, and differentially likely across the two mechanisms, the true welfare differences between them could be quite different from the observable revenue differences. To interpret our revenue differences in welfare terms, we must check for the presence of differences in the mechanism-specific 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, 45% of nominated parcels failed to receive a

[26] In appendix A.2, we report estimates of the models in Table 3 with the natural logarithm of bonus payments as the left-hand side variable.

[27] Table 1 shows imperfect balance between auction and negotiated leases across many of our conditioning variables, even after making the aforementioned sample restrictions. To verify that our results are still robust to any latent bias due to this imbalance, we estimate overlap-weighted treatment effects in Appendix A.4.

qualifying bid, so on a per-transaction basis, failure is quite common. The GLO often offers to sell these failed parcels again in future auctions, to the point that 72% 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 2. For each 10-mile by 10-mile grid, we compute the fraction of auction and negotiation parcels 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 4. The left-hand-side variable is a dummy indicating that at least one lease occurs during our sample period (2005-2016). Across different spatial controls, and accounting for differences in observables on the surface, these models reveal no statistically or economically significant differences in the probability of a parcel being leased in our sample. Given these results, it does not appear that any extensive margin corrections are necessary to interpret the bonus results reported above.

5.2 Robustness Checks

As we discussed in section 4, our key identifying assumption is that land that was initially owned by the state but sold between 1895 and 1931 is similarly valuable for today's hydrocarbon exploration as land from the same allocation that was sold after 1973 or is not yet sold. 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,

Figure 2: Time to First Lease for Auction and RAL Parcels

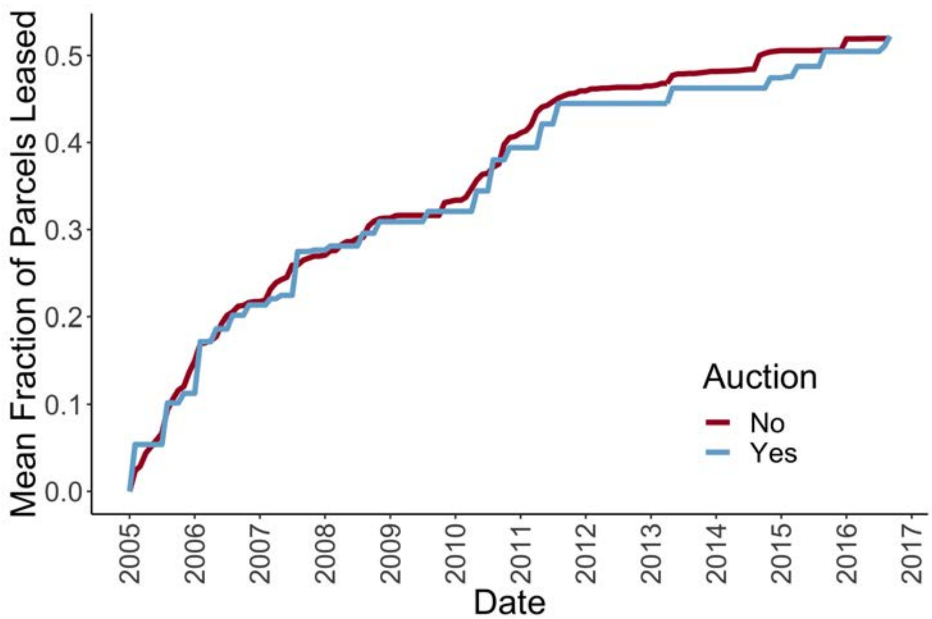

Average across 10 square mile grids of the fraction of parcels that have leased at least once since January 2005, by parcel type.

Table 4: Likelihood of Leasing and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.003 (0.020)	0.008 (0.019)	0.013 (0.016)	0.004 (0.019)	0.009 (0.018)	0.014 (0.016)
Grid	20	10	RF	20	10	RF
Extra	No	No	No	Yes	Yes	Yes
N	1,747	1,747	1,747	1,747	1,747	1,747
R^2	0.748	0.805	0.017	0.750	0.808	0.029

The dependent variable equals 1 if a parcel was ever leased and 0 otherwise. In columns 1-2 and 4-5, the size of the location bins, in miles, are indicated in the "Grid" row. Standard errors are clustered by grid in columns 1-2 and 4-5. Columns 3 and 6 use a Robinson (1988) transformation using random forests, with heteroskedasticity-consistent standard errors. All models include a spline in the size of the parcel in acres, and models with "extra" controls include shape regularity, surface cover, and distance to roads and water sources.

our empirical specifications include flexible spatial controls to account for any differences in geology across leases governed by the two mechanisms. Moreover, within narrowly defined locations, Table 2 shows that the land for auctioned and negotiated leases are similar on a host of observable characteristics. Finally, although Figure 2 shows leases on unleased parcels occur at similar times within these narrowly defined grids, our main specifications include additional space-by-time controls. This ensures that we are making comparisons between leases that transact in both similar places and similar times.

Nevertheless, if our identification assumptions were wrong, and auctioned leases were somehow easier to develop, then E&P's would rightly pay more for these parcels, but our conclusion that auctions have a causal effect on seller revenues would be incorrect. As an extreme test of this sort of reverse causality, we look for differences in bonus payments *conditional* on ex-post measures of lease value. We observe investment (in the form of drilling) and production (in the form of product-specific royalty revenues) for each lease, and can thus directly measure the extent to which accurately anticipated differences in these outcomes explain *ex ante* bonus payment differences. To do this, we add investment and production as covariates to our standard bonus payment regressions, as shown in Table 5. Columns 1 and 2 of this table show that while E&P companies are willing to pay more for parcels that are eventually drilled, they continue to pay more in auctions, even after conditioning on drilling decisions. The estimated coefficient on the Auction variable is still large, precisely estimated, and in the same ballpark as our main specification in Table 3. Columns 3 and 4 show that this remains true even if we condition on realized production. As in the first two columns, we do see evidence that leases which ultimate produce more output do earn higher bonus payments, but this does not diminish the difference in up front payments between auctioned and negotiated leases.

Another identification concern relates to a form of exclusion restriction implied by our empirical strategy. To conclude that the difference in bonus payments between auctioned and negotiated leases is caused by the mechanism itself, we must assume that there are no other channels through which the Relinquishment Act affects lease outcomes. One potential confounder is the difference in surface ownership between auctioned and negotiated leases. 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, relative to leases where the state controls the surface. If these constraints made negotiated leases more difficult to develop,

Table 5: Bonus Payments and Mechanism Type: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	700.39 (353.12)	704.57 (164.13)	681.51 (360.98)	681.95 (167.34)	630.53 (651.26)	858.58 (240.85)
Drilled	161.49 (207.09)	237.51 (117.96)				
Output			0.34 (0.34)	0.41 (0.15)		
Estimate	G10Y	RF	G10Y	RF	G10Y	RF
Sample	All	All	All	All	Private	Private
N	1,320	1,320	1,320	1,320	1,103	1,103
R^2	0.760	0.051	0.761	0.060	0.755	0.053

The dependent variable in each regression is bonus ($) per acre. In columns 5 and 6, the same is restricted to leases with strictly private surface ownership. Columns 1, 3, and 5 use fixed effects for year-by-10-mile grid, as well as quarter of sample. Columns 2, 4 and 6 use a random forest to semi-parametrically control for location and time. Standard errors clustered by location in columns 1, 3, and 5, and are heteroskedasticity-consistent in columns 2, 4 and 6. All models include a spline in acres, and linear terms in term length and royalty rate.

E&P companies would rationally pay less to lease them, but this difference in payment would not be caused by the difference in mechanisms.

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. Among this subsample, there are two types of auctioned leases. First, as previously discussed, a small number of leases on RAL parcels are auctioned when the state cannot locate the surface owner. Second, land sold from the PSF after 1973 has private surface ownership, just like leases on RAL parcels. 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 5 present our main bonus regressions specifications re-run on this restricted sample. Although the results in column 5 are less precise, due to the fact that it excludes nearly half of the auctioned leases, these estimates are remarkably consistent with the results in Table 3.[28] 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 pro-

[28]In Table A.3 we estimate all of the specifications in Table 3 using this restricted subsample. The point estimates in this table are generally higher than those in 3, and many are precisely estimated.

visions 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.

6 Allocative Efficiency Results

Our results on the causal effect of auctions on bonus payments are consistent with predictions from a strand of the theory literature that the "competitive effects" of auctions generate tangible benefits for sellers (Bulow and Klemperer, 1996, 2009). This literature also suggests that auctions should generate higher total welfare than other less formal mechanisms, by reliably awarding goods to the buyers who values them the most. For example, Klemperer (2004) argues that "a well-designed auction is the method most likely to allocate resources to those who can use them most valuably." Motivated by these additional predictions from the theory literature, we use our drilling and production data to measure the causal effect of auctions on allocative efficiency.

In order for auctions to produce better allocative outcomes than negotiations, two conditions must be true. First, it must be the case that potential lessees are heterogenous in their cost of investment or ability to produce output. 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. We know these differences in costs or productivity must exist because 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 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. And 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.

Heterogeneity across potential lessees won't necessarily generate allocative differences across the two mechanisms if negotiations somehow select the same winning firms as auctions do. In addition to characterizing potential sources of heterogeneity among firms, we can also show that auctions and negotiations generate noticeably different allocations of firms to leases. We do this by tabulating auction and negotiation "market shares" for each of the ten most active lessees, as shown in Table 6.[29] 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 ($p < 2 \times 10^{-16}$).[30]

Because E&P companies often focus their leasing activities in a specific area, and because different areas have different shares of auctioned and negotiated leases, these differences in market shares across the mechanism types may simply reflect differences in the distribution of a firm's "interest" across basins. To verify that this kind of heterogeneity is not driving these differences in market shares across mechanism types, we also replicate this exercise within leases overlying the two largest shale basins in Texas, the Permian and the Eagle Ford. We can similarly reject a null hypothesis of equal proportions for the top 10 most active lessees in each basin. Thus, the data on firm assignment is consistent with the idea that the two mechanism types generate different allocations. One major consequence of this difference is differential concentration across mechanisms. Table 6 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.[31]

[29] It is important to emphasize that firm identities are observed with error in our data. Additional information on our process for cleaning names is provided in Appendix B.

[30] Chi-squared tests of equal proportions for the top 20 and 40 most active lesses are similarly rejected.

[31] The gap between the total share of top 20 auction winners and negotiation winners is about 11%, and the gap for the top 40 is about 8%.

Table 6: Top 10 Auction Winners and Negotiators

Firm	Leases	Auction Share	Negotiation Share
CHESAPEAKE	114	0.183	0.035
LEWIS OPERATING	83	0.022	0.085
ENERGEN	80	0.061	0.060
PETROHAWK	72	0.085	0.038
PETRO HUNT	69	0.007	0.077
CIMAREX	59	0.039	0.048
ANADARKO	56	0.048	0.040
DEVON	33	0.059	0.007
BP PRODUCTIONS	31	0.000	0.036
RANGE PRODUCTION	31	0.043	0.013
ALL OTHERS	692	0.454	0.562

While it is not hard to imagine that firms will differ in their use value for any given parcel, and the data is supportive of the idea that auctions allocate differently than negotiations, actually testing whether auctions select *better* lessees is complicated by the fact that we don't directly observe the latent productivity for any lease-lessee pair (and we don't even observe who the set of potential lessees are for the negotiated leases). However, under any source of heterogeneity between firms, vertical or horizontal, if auctions have higher allocative efficiency, we should expect to see that auctioned leases are put to more productive use than negotiated leases. This is something that we can directly measure using our royalty payment data.

Output from the oil and gas industry is administratively recorded and reported at high frequency. This output is undifferentiated (aside from oil vs. gas) and the market for it is competitive, so direct output comparisons have a clear normative interpretation. The primary challenge in using production data implied by product-specific royalty payments is that leases generate output over many years, and all producing leases in our sample are right censored. To deal with this, we begin by estimating the impact of auctions on the probability that a lease is ever drilled. We restrict the sample to leases whose primary term has concluded by the end of June, 2018. Within this group, drilling is an important indicator of the eventual value of a lease, yet the outcome is unaffected by right censoring.

The top panel of Table 7 presents the results. The model specifications in each column are identical to those in Table 3, showing the effects of mechanism type on lease outcomes, under various spatial and temporal controls, as well as the inclusion of "extra" covariates. Across specifications, auctioned leases are 8 to 18 percentage points more likely to be drilled, and in half of these specifications, the estimates are statistically different from zero. As just

38% of negotiated leases are ever drilled, this increase in the probability of drilling is large in relative terms. Although the models in columns 2, 4 and 6 are imprecisely estimated, the general picture across the specifications suggests large differences in the probability of investment, providing our first set of evidence linking a transaction's mechanism type to its real outcomes.

Table 7: Likelihood of Drilling, Lease Output, Lease Revenue and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)
Auction - Drilling	0.116	0.082	0.179	0.096	0.132	0.080
	(0.044)	(0.050)	(0.078)	(0.061)	(0.042)	(0.051)
R^2	0.397	0.657	0.747	0.500	0.038	0.666
Auction - Output	117.48	85.05	114.18	116.44	141.28	96.19
	(37.97)	(40.30)	(66.77)	(38.53)	(41.53)	(42.46)
R^2	0.557	0.783	0.859	0.630	0.021	0.789
Auction - Revenue	1358.10	880.08	1210.34	1284.82	1927.48	922.79
	(454.11)	(578.28)	(838.74)	(540.47)	(455.34)	(586.90)
R^2	0.592	0.791	0.890	0.645	0.027	0.796
Grid	10	10	10	20	RF	10
Time	Q	GY,Q	GYQ	GY,Q	RF	GY,Q
Extra	No	No	No	No	No	Yes
N	1,109	1,109	1,109	1,109	1,109	1,109

The dependent variables are an indicator for whether the lease is associated with royalty revenue in sample (Drilling), discounted barrels of oil equivalent per acre (Output), and total revenue per acre, which is the sum of the bonus and discounted royalty payments (Revenue). 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 Robinson (1988) transformation using random forests, with heteroskedasticity-consistent standard errors. 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, and distance to roads and water sources. The sample includes all leases whose primary term ends before June, 2018. Negotiated leases are drilled 38% of the time. The average negotiated lease generates 192 bbl of discounted BOE per acre, and $2,575 of total seller revenues per acre.

Next, we estimate the impact of mechanism type on lease output, measured in discounted barrels of oil equivalent per acre, which is the actual object of import for the allocative efficiency question. Before discussing the results, it is worth noting that this variable is incredibly skewed: for leases that are drilled in our sample, the difference between the 10^{th} and 90^{th} percentiles of output spans more than three orders of magnitude. A natural solution to this right skewness would be to examine differences in output across leases in relative terms, by using the natural logarithm of output as the dependent variable. However, as described above, fewer than half of leases are ever drilled, and as such generate zero output in the real sense (i.e., this is not just a selection problem). In this situation, adding a small constant to these zeros to facilitate the logarithmic transformation is unlikely to be innocuous. We thus proceed in levels, acknowledging that the skewness of the data will make these measurements inherently imprecise.

The middle panel of Table 7 presents the results. In all specifications, auctioned leases produce substantially more than negotiated leases, and all estimates are statistically different from zero. Even at the low end of these estimates (column 2), the difference of roughly 85 barrels of oil equivalent per acre is 44% of the average output on negotiated leases. Using alternative spatial and temporal controls, the differences between auctioned and negotiated leases are even larger. For example, the point estimate in column 5 is 74% of the average output of negotiated leases.

This increased output implies that not only are auctions more efficient, they likely have an even larger impact on seller revenues than the bonus regressions imply. The bottom panel of Table 7 measures this formally, using the sum of bonus payments and discounted royalty revenues as the dependent variable. Across all six specifications, we find large differences in total revenue, though the differences for specifications with finer spatial controls are imprecise due to the small sample and skewness of royalty payments.

6.1 Unpacking the source of heterogeneity

Table 7 provides evidence that auctions allocate leases to firms who are more likely to drill them, and who produce more output with them, a result that is consistent with some of the theory literature's predictions of allocative efficiency gains from auctions. In this section, we provide statistical evidence regarding the relative contribution of vertical or horizontal productivity differences between firms in generating these results. A key 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 simply make landowners aware of *who* productive lessors are and nudge lessors to select them. Conversely, 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 discussed above, vertical productivity factors generate persistent differences across firms. By definition, these factors are constant across leases within a firm. Thus, if productivity differences between auctions and negotiations are driven primarily by vertical differences in the firms these mechanism allocate to, then comparisons of drilling and output *within a firm* should reveal no auction treatment effect. We test this hypothesis in Table 8, which reports estimates of our main drilling and output regressions, with and without fixed effects for the identity of the lease winner. Even after conditioning on firm identity, bonus payments, the likelihood of drilling, and lease output are all still larger, by a similar magnitude, in auction leases relative to negotiated leases. If variation in allocation across vertically differentiated firms was an important driver of the difference in lease outcomes across the two mechanisms, we would expect to see that the differences in outcomes, conditional on firm identity, would be closer to 0 than the unconditional differences. If anything, the opposite seems to be true, although confidence intervals overlap considerably. We also perform this exercise with bonus as the dependent variable, finding similar results: auctions pay more, even within firm.

Table 8: Effects of Firm Composition and Mechanism Type on Lease Outcomes

	Bonus	Bonus	Drilled	Drilled	Output	Output
Auction	708.85 (342.20)	719.36 (437.69)	0.082 (0.050)	0.177 (0.090)	85.05 (40.30)	175.00 (84.29)
Firm FE	No	Yes	No	Yes	No	Yes
N	1,320	1,320	1,109	1,109	1,109	1,109
R^2	0.760	0.825	0.657	0.748	0.783	0.859

The dependent variable is bonus per acre (columns 1 and 2), a dummy variable for whether the lease is drilled (columns 3 and 4) or discounted output in BOE terms per acre (columns 5 and 6). 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 lessee-lease shocks as a determinant of differences between auctioned and negotiated leases? Unfortunately, we are not aware of a direct test for this hypothesis. However, 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 or more likely to drill than others, we'd expect to see a consistent ranking of auction 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 more than ten times. Table 9 lists these pairs and tabulates the probability that the alphabetically earlier firm (Firm A) bids higher than the later firm (Firm B). If firm-lease match was less important than vertical differences between firms, we'd expect to see that one firm consistently bids higher than the other. What we observe is the exact opposite: for 8 of the 9 pairs, the fraction of the time that one firm wins more than the other is statistically identical to a coin toss.

Table 9: Bid ranking for top auction pairs

Firm A	Firm B	Auctions	Share A > B
CIMAREX	ENERGEN	31	0.52
CIMAREX	CONOCO PHILLIPS	19	0.79
CONOCO PHILLIPS	ENERGEN	19	0.37
CIMAREX	RESOLUTE	19	0.53
ENERGEN	RESOLUTE	19	0.42
COG	RANGE PRODUCTION	17	0.41
CONOCO PHILLIPS	RESOLUTE	17	0.53
CIMAREX	MARSHFIELD OIL AND GAS	12	0.67
ENERGEN	MARSHFIELD OIL AND GAS	12	0.67

7 Discussion

Texas mineral leases allocated by auctions generate more revenue for mineral rights owners and are better matched to firms who can use these minerals productively, relative to leases allocated by informal, decentralized negotiations. What features of the unobserved "RAL mechanism" underlying these negotiated transactions are responsible for these differences? A significant limitation in answering this question is that 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 outcome differences fit within existing mechanism comparisons considered by the literature.

Section 6 shows that negotiations do not always allocate leases to the highest value user. One possible explanation for this phenomena is that the set of potential lessees for negotiations and auctions are the same, but negotiations choose among these lessees in a sequential fashion, while the auctioneer selects among all participants simultaneously.

The theory literature offers conflicting opinions about the efficiency costs and benefits of a sequential process relative to a simultaneous process. Bulow and Klemperer (2009) show that sequential mechanisms *can* perform worse than an auction, both in terms of generating seller revenues and allocative efficiency, if a "good enough" bidder randomly arrives early enough in the sequential process. However, this possibility 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 (up to the point when a transaction occurs), the improved performance of auctions in our setting suggests that entry choices by E&P companies are not especially "selected," as in the Roberts and Sweeting (2013) mechanism.

An alternative explanation consistent with the decline in allocative efficiency is that RAL leases actually transact using an auction, but "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, the possibility that RAL surface owners are unable (or unwilling) to acquire offers from all potential lessees seems more than plausible. 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 search to determine who actually owns it. It is thus likely that informal mechanisms for the latter 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.

While these two mechanisms could rationalize our empirical results, it is important to note that neither 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 any 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.[32] 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, nowadays 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 is a 67% increase in bonus price, this gain from using an auction appears to far exceed the cost.[33] 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.[34]

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 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 9, we tabulate the number of auctions with 1, 2, 3, 4 or 5+ bidders, and within those groups, compute the average bonus payment per acre and the median reserve margin. More than two-thirds 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).

[32] Survey 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].

[33] Note 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.

[34] Indeed, 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.

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 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 67% larger up front payments, and that auctions produce 44% more 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).

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.

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.

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.

Klemperer, P., *Auctions: Theory and Practice* Auctions: Theory and Practice, Princeton University Press, 2004.

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 2014.

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.

Milgrom, Paul, *Putting Auction Theory to Work*, Cambridge University Press, January 2004.

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, "Optimal auction design," *Mathematics of operations research*, 1981, *6* (1), 58–73.

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.

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

Wang, Zhe, "Designing M&A Selling Mechanisms: Go-Shop Negotiations," *Available at SSRN 3287321*, 2018.

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

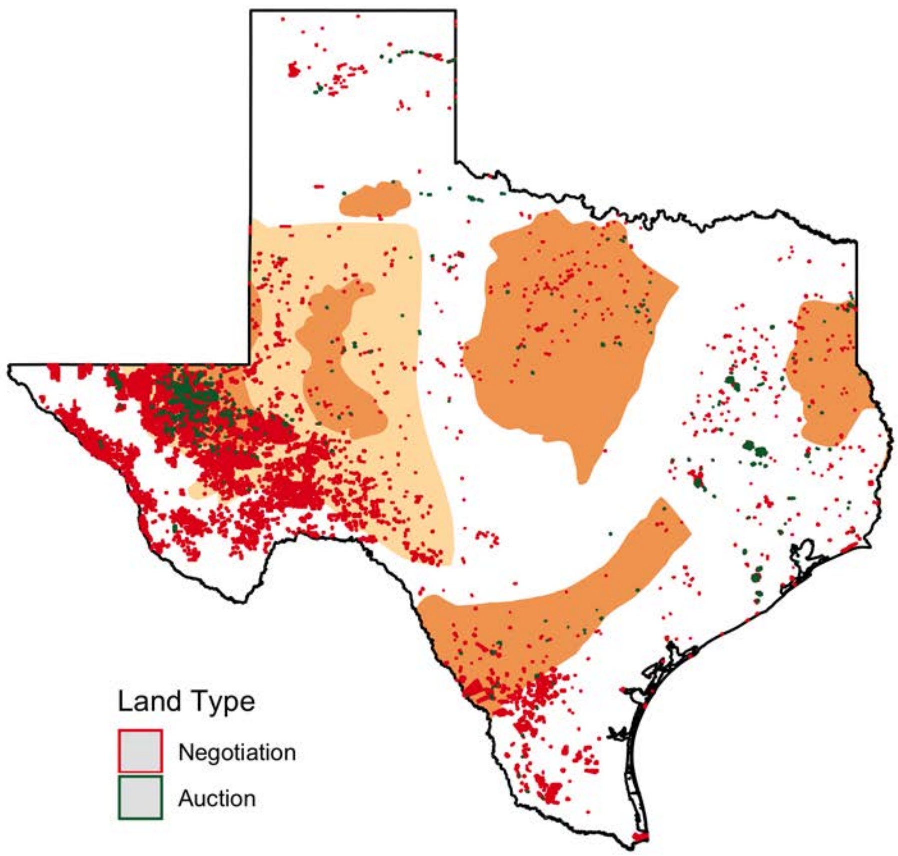

Figure A.2: Example of Sample Lease Type Overlap

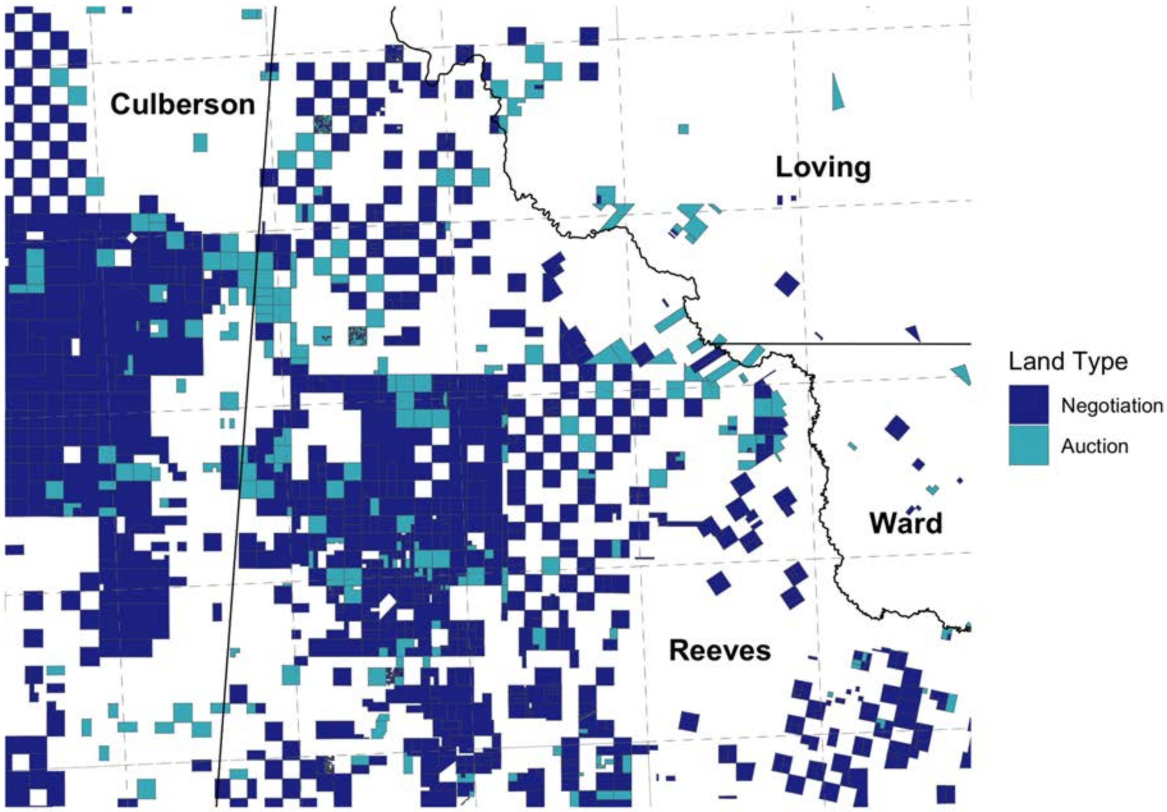

A.2 Log Bonus Results

Table A.1: log(Bonus) Payments and Mechanism Type

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.35	0.40	0.40	0.41	0.43	0.39
	(0.08)	(0.10)	(0.13)	(0.10)	(0.06)	(0.10)
Grid	10	10	10	20	RF	10
Time	Q	GY,Q	GYQ	GY,Q	RF	GY,Q
Extra	No	No	No	No	No	Yes
N	1,320	1,320	1,320	1,320	1,320	1,320
R^2	0.919	0.970	0.983	0.954	0.730	0.970

The dependent variable in each regression is the natural logarithm of the lease's bonus payment. 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 Robinson (1988) transformation using random forests, with heteroskedasticity-consistent standard errors. All models include a spline in acres and linear terms for 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.

Table A.2: log(Bonus) Payments and Mechanism Type: Robustness

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	0.39	0.42	0.39	0.41	0.40	0.48
	(0.10)	(0.05)	(0.10)	(0.06)	(0.13)	(0.08)
Drilled	0.12	0.27				
	(0.06)	(0.05)				
Output			0.00	0.00		
			(0.00)	(0.00)		
Estimate	G10Y	RF	G10Y	RF	G10Y	RF
Sample	All	All	All	All	Private	Private
N	1,320	1,320	1,320	1,320	1,103	1,103
R^2	0.970	0.738	0.970	0.741	0.972	0.713

The dependent variable in each regression is the natural logarithm of the lease's bonus payment. Columns 1, 3, and 5 use fixed effects for year by 10-mile grid, as well as quarter of sample. Columns 2, 4 and 6 use a random forest to semi-parametrically control for location and time. Standard errors clustered by location in columns 1, 3, and 5, and are heteroskedasticity-consistent in columns 2, 4 and 6. All models include a spline in acres, and linear terms in term length and royalty rate.

A.3 Private Surface Owner Models

Table A.3: Bonus Payments and Mechanism Type: Private Surface Owners Only

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	633.09 (342.46)	630.53 (651.26)	1399.95 (1030.94)	783.22 (336.46)	858.58 (240.85)	653.24 (626.26)
Grid	10	10	10	20	RF	10
Time	Q	GY,Q	GYQ	GY,Q	RF	GY,Q
Extra	No	No	No	No	No	Yes
N	1,103	1,103	1,103	1,103	1,103	1,103
R^2	0.566	0.755	0.907	0.665	0.053	0.759

The dependent variable in each regression is the lease's bonus payment. 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 Robinson (1988) transformation using random forests, with heteroskedasticity-consistent standard errors. All models include a spline in acres and linear terms for 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.

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.4 using the "causal forest" estimator from the `grf` R package, as described in Athey et al. (2019). Our results are close to those presented in Tables 3 and 7.

Table A.4: Overlap-weighted ATE Estimates

	Bonus	Bonus	Drilled	Drilled	Output	Output
Auction	841.27 (141.09)	821.48 (137.78)	0.100 (0.041)	0.096 (0.039)	170.37 (44.58)	183.82 (43.68)
Extra	No	Yes	No	Yes	No	Yes
N	1,595	1,595	1,328	1,328	1,328	1,328

Overlap weighted average treatment effect estimates for Bonus per acre (columns 1 and 2), Drilled (columns 3 and 4) and Output per acre (columns 5 and 6). Columns 4 through 6 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, and distance to roads and water sources for the "extra" columns 2, 4 and 6.

A.5 Auction Statistics

Table A.5: Auction Results by Number of Bidders

Bids	Auctions	Fraction	Bonus ($/Acre)	Bid/Reserve (med.)
1	263	0.57	1,418	1.66
2	88	0.19	1,877	2.08
3	35	0.08	3,014	3.61
4	32	0.07	2,795	3.26
5	22	0.05	6,225	3.61
5 +	19	0.04	6,027	4.71

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

Appendix B Data Cleaning

B.1 Sample construction

Table A.6: Sample Construction

	Drop Reason	Negotiation	Auction
All Leases		4,012	915
	Not on Shale	2,359	550
	Missing Value	2,314	550
	Less Than 10 or Greater Than 1,000 Acres	1,869	506
	Gross and Net Acreage Differ	1,406	504
	Undivided Interest	893	477
	Term Less Than 1 Year	883	477
	Cancelled or Withdrawn	874	474
	Negotiated State Lease	861	474
	Lessee Owns RAL Surface	860	474
	PSF Acquired Land	860	460
Final Sample		860	460

Additional discussion provided in section 3.1.

B.2 Firm Names

Though we observe the name of the firm on the lease, E&P companies sometimes use intermediaries 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.7, 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 pleases pay about $700 more per acre 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.7: Bonus Payments and Mechanism Type: Robustness to RAL Lease Addenda

	(1)	(2)	(3)	(4)	(5)	(6)
Auction	708.85	717.21	822.94	845.30	667.89	734.65
	(342.20)	(160.32)	(284.62)	(134.66)	(327.38)	(192.86)
Pages			-16.80	-14.19		
			(45.91)	(13.03)		
Clauses			29.69	37.55		
			(24.17)	(13.20)		
Surface Protection					-2.93	-122.36
					(82.63)	(88.66)
Payment Terms					-117.63	-18.60
					(98.63)	(73.46)
Location Requirements					-117.73	-18.60
					(155.67)	(113.06)
Pugh Clause					-135.40	63.13
					(229.21)	(105.04)
Cleanup Terms					-173.46	-11.97
					(171.19)	(147.73)
Livestock Protection					-76.94	-135.92
					(287.15)	(166.70)
On-site Water Use					342.62	351.57
					(396.32)	(263.00)
Waste Management					36.45	-326.05
					(163.81)	(205.03)
Definitional Changes					43.03	64.54
					(70.45)	(71.50)
Pollution Protection					-116.19	264.17
					(179.74)	(184.80)
Infrastructure Constraints					-20.44	-174.42
					(197.91)	(162.97)
Caliche Use					-698.72	191.84
					(757.15)	(359.97)
Additional Fees					14.07	4.77
					(147.35)	(136.93)
Time Constraints					296.93	120.46
					(260.56)	(276.24)
Miscellaneous					123.21	81.14
					(191.89)	(162.51)
Grid	10	RF	10	RF	10	RF
Time	GY,Q	RF	GY,Q	RF	GY,Q	RF
N	1,320	1,320	1,320	1,320	1,320	1,320
R^2	0.760	0.046	0.762	0.063	0.763	0.059

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 random forest in lease latitude, longitude and effective date.