Multidimensional Auctions of Option Values

Yunmi Kong[†] Isabelle Perrigne[‡] Quang Vuong[§]

September 27, 2019

Abstract

This paper conducts a structural analysis of auctions of contracts when private information is multidimensional. These involve a trade-off between adverse selection and moral hazard. Specifically, we study oil lease auctions where bids consist of both a cash payment and a royalty rate on extraction revenue. Upon modeling the lease as an option - as the winner chooses whether to exercise it ex-post - we derive optimal bids under an unspecified scoring rule. We nonparametrically identify bidders' types from their bids and develop a nonparametric estimation procedure to recover their joint distribution. Analyzing cash-royalty auctions in Louisiana, our rich model allows us to compare their performance to cash auctions with fixed royalty in terms of moral hazard, allocative effects, information rents, social surplus, and government revenue. We find that endogenous royalties exacerbate adverse selection and moral hazard; fixing royalties is revenue-superior to allowing it as a bid component. In addition, we assess the effects of changing the lease duration and exploiting fluctuations in oil prices.

We would like to thank the Louisiana Department of Natural Resources for providing the data. We also thank seminar/conference participants at WUSTL Olin, Texas A&M, the Lancaster conference on auctions, IAAE, SITE, and the Banff empirical microeconomics workshop.

[†]Rice University, Department of Economics. E-mail: yunmi.kong.01@gmail.com
[‡]Rice University, Department of Economics. E-mail: iperrigne@gmail.com
[§]New York University, Department of Economics. E-mail: qvuong@nyu.edu

1 Introduction

When a principal wishes to contract with an agent, there may be many potential agents from which to choose, who each have different types that the principal cannot observe. The principal must decide how to choose among these agents and set contract terms. In designing such a mechanism, he faces a trade-off between adverse selection and moral hazard. A mechanism that links the principal's payoff more closely to outcomes - e.g. via a sharing rule - reduces agents' information rents, but it also reduces agents' incentives to exert effort. In this context, one mechanism proposed by theory is an auction of contracts in which the auction endogenously determines both an upfront payment (cash) and a sharing rule (royalty). Specifically, McAfee and McMillan (1987) and Laffont and Tirole (1987) explain that auctions of this form can achieve an optimal trade-off between information rents yielded to agents and moral hazard, when agents have one-dimensional types.

Indeed, formal and informal auctions which determine both a cash payment and royalty are commonly observed; examples include government-firm, author-publisher, patent holder-licensee, entrepreneur-acquirer, landlord-sharecropper, and sports association -broadcaster contracting, among others. However, whether these auctions perform well in practice is an empirical question. First, the optimal auction proposed in theory involves the principal designing a menu of cash-royalty contracts; real-world auctions may not be implemented in this way. Second, agents may have multidimensional types, while existing optimality results concern one-dimensional types. Cash-royalty auctions have the potential to perform better than simpler auctions - e.g. which fix the royalty and have bidding on the cash component only - but they may in fact do worse depending on how they are implemented.

This paper conducts a structural analysis of auctions of contracts with multidimensional private information. Specifically, we study government oil lease auctions where bids consist of both a cash payment and a royalty rate on extraction revenue. Each bidder chooses his two-dimensional bid based on his two-dimensional type. Meanwhile, the auctioned leases grant the right, but not the obligation, to develop a tract of land for oil production. As a result, the royalty component of the bid induces moral hazard by reducing the winning firm's incentive to develop the tract. The fact that there is no obligation is of first-order importance empirically; the majority of publicly leased tracts are left undeveloped. Therefore, the contract takes the form of a real option in our model. This endogenizes a bidder's value for the contract to explicitly account for the probability that he will not develop the tract and for how the royalty component of his bid affects this probability. Royalties chosen endogenously by bidders also have important implications for adverse selection. While higher royalty levels generally reduce firms' information rents,[1] we show that endogenous royalties exacerbate adverse selection because the firms most willing to bid up royalties are less desirable types.

The data we use are from Louisiana's cash-royalty auctions of oil leases on state lands. Though bids are two-dimensional, there is no specified scoring rule, so we cannot rely on properties of specific scoring functions in our analysis. However, much can be learned - by bidders and by the econometrician - from observations of the state's choice over many auctions. In particular, we document that the state favors both higher cash payments and higher royalties, in contrast to some optimal mechanisms proposed in theory. Another pattern we see in the data is that the within-auction distribution of bids is scattered across two dimensions, in a manner strongly suggestive of multidimensional rather than one-dimensional bidder types.

As mentioned previously, bidders have two-dimensional types (θ_1, θ_2) in our model. The first component, θ_1, represents productivity or expected production volume, and the second component, θ_2, represents cost. From the perspective of a bidder who has bid royalty rate a, the contract is essentially an option to obtain $(1 - a)\theta_1$ expected units of a commodity at a cost of θ_2. Whether the bidder will exercise this option depends on the future price path of the commodity. We value this contract by adapting insights from the option pricing literature which exploit properties of geometric Brownian motion. Building these option values into an auction model, we derive optimal cash-royalty bids as a function of bidders' types under an unspecified scoring rule. As part of this analysis, we show that bidding a higher royalty percentage is less costly for undesirable "weak" types - those with low productivity and/or high cost - than for "strong" types. This has important implications for adverse selection; the royalty component of bidding effectively provides weak types a cheaper currency with which to bid. The insight is similar in spirit to those offered by Che and Kim (2010) and Skrzypacz (2013) section 5.1.

Given this model, we nonparametrically identify bidders' types from their bids. Specifically, the model yields a system of two first-order conditions that characterizes

[1]See DeMarzo et al. (2005).

the bidder's choice of bid as a function of his type. As the bid is observed, this is a system of two equations in which the bidder's type components are the two unknowns. The mathematical intuition for identification is that this system of equations has a unique solution. Economically, the separate identification of θ_1 and θ_2 is related to the fact that bidders' types cannot be reduced to scalar representations in our model. Royalties are levied on revenue not profit, so the distinctive roles played by productivity versus cost in determining the bid cannot be reduced to or represented by the bidder's overall value for the contract. Using observed bids, we recover types for any bidder who submits a bid, regardless of whether he won or whether the lease was developed.

Based on the identification argument, we develop a nonparametric estimation procedure to recover the joint distribution of type components. As the probability of winning plays an important role in bidders' choice of bid, the most important step of estimation is estimating the probability of winning as a function of the cash-royalty bid. We exploit structure to estimate this probability well in a finite sample: first, we estimate via kernels the bivariate bid distribution; second, we estimate via sieves the probability that the state favors bid A over bid B, as a function of the bid components of A and B. Then, the probability of winning is estimated as their composite: the probability that a given bid is favored over competing bids drawn from the observed bid distribution. After this step, we follow the constructive identification argument to estimate bidders' two-dimensional types.

Using our estimated structural model, we compare Louisiana's cash-royalty auction to counterfactual fixed-royalty auctions, in which the principal fixes a common royalty rate, and bidders bid only on the cash payment amount. The model is sufficiently rich to allow comparison in the particulars of moral hazard, allocative effects, and information rents, the combined effects of which lead to an overall comparison of social surplus and government revenue. We find that the cash-royalty auction in Louisiana disproportionately weakens cash competition relative to the royalty level it achieves. As a consequence, total government revenue - the sum of cash payments and royalties - is lower than in fixed-royalty auctions, even when the fixed royalty is set at conventional rates rather than chosen optimally. The fixed-royalty auction also outperforms counterfactual scoring auctions, in an important departure from properties derived for multi-attribute auctions with no moral hazard. Finally, we assess policy instruments beyond auction design, such as changing the duration of the lease and exploiting fluctuations in oil prices. Higher oil prices have a positive impact on government revenue that is much more than proportional to the price increase, while increasing the lease duration does not increase the probability of option exercise.

Related Literature In studying multidimensional auctions of contracts, this paper primarily relates to two streams of literature: the multi-attribute auctions literature and the auctions of contracts literature. In addition, this paper is related to work on real options and empirical oil lease design. To our knowledge, this is the first paper to build multidimensional private information and bidding into auctions of contracts, which are subject to adverse selection and moral hazard.

The cash-royalty auctions we study have some affinity to price-quality procurement auctions, in which the winner is chosen based on the quality or speed of her work as well the price she names for it. An important difference is that there is no moral hazard in price-quality auctions. In this literature, Yoganarasimhan (2016) and Krasnokutskaya, Song, and Tang (2018) share a similarity with our paper in that the scoring rule is unknown to the econometrician, but they differ from ours in that the non-cash component of the bid is considered an exogenous attribute of the bidder.[2] In the absence of a scoring rule, Yoganarasimhan (2016) estimates probabilities of winning as a function of all the bidders' prices and qualities. In Krasnokutskaya et al. (2018), the quality component of the bid is also unobserved, so the paper first estimates differences in quality and then estimates the linear preferences which determine the winner. Meanwhile, Che (1993), Asker and Cantillon (2008, 2010), and Lewis and Bajari (2011) study quasi-linear scoring rules when both bid components are endogenous, while Takahashi (2018), Sant'Anna (2018), and Hanazono, Hirose, Nakabayashi, and Tsuruoka (2016) study price-quality ratio, interdependent, and general scoring rules, respectively.

Auctions of contracts motivated in a principal-agent setting have been studied by a number of papers in the theory literature. In addition to those we reference in the introduction, Hansen (1985) and DeMarzo, Kremer, and Skrzypacz (2005) explain that linking the principal's payoff to the agent's output through contingent payments or royalties reduces bidders' information rents. Specifically considering contracts that take the form of a real option, Board (2007) derives the principal's optimal

[2] Nakabayashi (2013), Laffont et al. (2018), and Andreyanov (2018) are other examples of multi-attribute auctions in which the non-cash component of the bid is exogenous.

mechanism, and Cong (2019) derives the optimal design among standard security bids. In both cases, the optimum involves a menu of upfront cash and contingent payments – either a lump-sum or royalty to be paid upon option exercise – that are negatively correlated. Skrzypacz (2013) surveys the related literature. Importantly, this literature is concerned with one-dimensional agent types; in general, the literature on multidimensional types is sparse due to the difficulties associated with a lack of exogenous ordering of the type space.

This paper also relates to Bhattacharya, Ordin, and Roberts (2018) and Herrnstadt, Kellogg, and Lewis (2019), which study empirical oil lease design through the lens of real options. Specifically, Bhattacharya et al. (2018) analyze data from fixed-royalty auctions and the timing of option exercise to study questions of auction design, and Herrnstadt et al. (2019) examine lease data to study how royalties and lease expiration deadlines help mitigate information rents and moral hazard. Our paper contributes to this literature by analyzing an auction in which the royalty is a bid component as well as cash and by analyzing bidders with multidimensional types.

2 Data

2.1 Introduction

The Louisiana Department of Natural Resources (DNR) sells oil leases on lands owned by the state of Louisiana and its agencies. We study data from these sales that occurred for onshore[3] leases of at least 5 acres between 1974 and 2003. A lease grants the lessee the right, but not the obligation, to develop the leased tract of land for oil production. The fact that there is no obligation to develop is a key feature of these leases; the Department of the Interior reports that as of the end of 2011, "approximately 56 percent of total acres of public land under lease in the Lower 48 States [..] are not undergoing either production nor exploration activities."[4] For our analysis, we restrict our attention to leases for which we could identify the township(s) of location.[5]

[3]With tract kind codes 2 and 4 in the Louisiana DNR categorization.

[4]"Oil and Gas Lease Utilization, Onshore and Offshore: Updated Report to the President," U.S. Department of the Interior, May 2012.

[5]A township is a 6-by-6 mile square in the Public Land Survey System.

2.2 Bids

To bid for a tract of land that Louisiana has advertised for leasing, a firm must fill out a bid form specifying the terms it offers for the lease. While a firm may specify that it wants to lease an arbitrary portion of the advertised tract, we will only analyze auctions where all bids are bidding for the entire tract, each bidder name is associated with one bid, and it is clear which bid won. This ensures that all bidders in an auction are bidding on the same piece of land and lets us focus our analysis on contract terms. The bid form asks the bidder to specify a cash payment, a royalty rate, the duration of the lease (up to a maximum of 3 years), and any additional considerations the bidder wishes to offer. The first two components will emerge as the deciding factors; we discuss each component in turn.

The cash payment is an amount that the winning bidder pays to the state even if he ultimately does nothing with the leased tract; it is non-contingent. It is the sum of two parts: an immediate payment and a "rental" which is paid annually during the life of the lease thereafter. The state stipulates that the rental must equal at least half of the immediate payment; in 94% of bids in our sample, the rental equals half of the immediate payment. We define the cash component of the bid as the immediate payment plus the discounted present value of rentals. Next, the royalty is a percentage of production revenue that the lessee pays to the state, contingent on producing. Production revenue equals production volume times the price obtained for oil. Note that the royalty is levied on revenue, not on profit. This is the norm for mineral leasing across the U.S., likely due to the administrative burden that verifying costs would entail. The minimum acceptable royalty in our sample is 12.5%.

The duration of the lease determines the number of years for which the lessee holds the right to develop the leased tract. If there is no production during that time, the lease expires; if there is production, the lease remains in effect for as long as production continues. For the onshore leases in our sample, Louisiana limits the duration to a maximum of 3 years. Though bidders are free to specify a shorter duration, 98% of bids are for 3 years in practice. We restrict our sample to auctions with such bids so that lease durations are homogeneous. Finally, bidders may offer additional consideration, e.g. an obligation to drill a well within a certain period of time, but none of the bids in our sample offer any.

Therefore, the bids in the auctions we study are effectively two-dimensional; the contracts bid differ in cash payment and in royalty rate. We refer to this two-

Figure 1: Scatterplot of cash-royalty bids, onshore 1974-2003

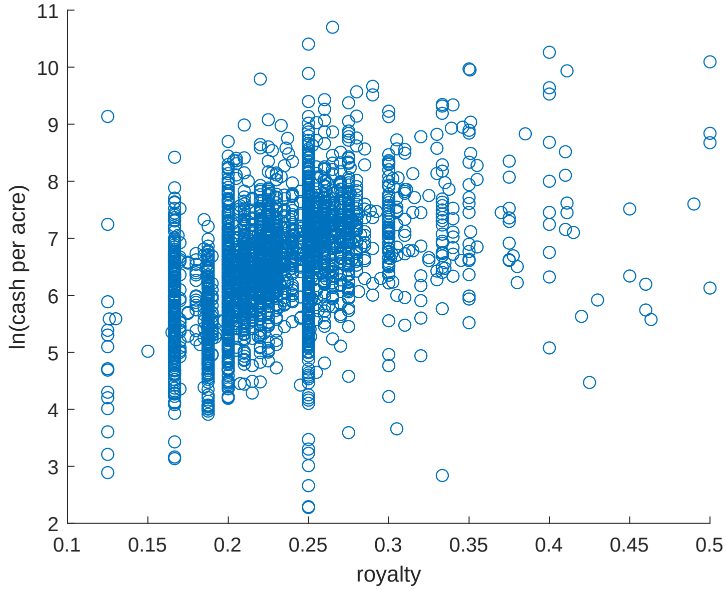

dimensional auction as the cash-royalty auction. This auction format is an alternative to the one-dimensional auction used by the federal Bureau of Land Management and some other states such as New Mexico, in which the government fixes a common royalty rate, and bidders bid only on the cash payment amount. We refer to the latter auction format as the fixed-royalty auction. Unlike the fixed-royalty auction, the cash-royalty auction allows each bidder to bid his own choice of royalty rate flexibly and endogenously.

Figure 1 displays a scatterplot of the cash-royalty bids observed in our data. Dollar amounts are deflated by the GDP implicit price deflator and stated in 2009 dollars. The median and mean cash bid are \$655 and \$1,020 per acre, respectively; the cash distribution is skewed and is closer to a log-normal than a normal distribution. The median and mean royalty bid are both 23%. As a comparison, the prevailing royalty rate on privately held land – owned by small landowners, farmers, etc. – is 25%. While there are outliers, royalty bids are concentrated between 15% and 30%. Meanwhile, the cash and royalty components of a bid are positively correlated; the raw correlation coefficient between the log of the cash bid per acre and the royalty bid is 0.48.

Table 1: Bidding and winning in two-bidder auctions

There is a dominant bid	62%
Dominant bid won	99%
Dominant bid lost	1%
There is no dominant bid	38%
Higher cash won	60%
Higher royalty won	40%

2.3 Which bid is chosen?

Unlike in a one-dimensional auction, what constitutes a winning bid is not immediately obvious in a multi-attribute auction. As such, some multi-attribute auctions specify a "scoring rule," which is a function used to convert a multi-attribute bid into a scalar score for purposes of determining a winner. Other multi-attribute auctions, including this Louisiana auction, do not specify a scoring rule. Here, technical staff make recommendations to the State Mineral and Energy Board regarding which bid is most advantageous to the state, and the Board makes a final decision as to the winner. Our inquiry into the finer details of the process yielded a response that the "geological and engineering staff look at each bid and make the determination [...] by looking at all factors involved." A few bidders who replied to our inquiries confirmed the absence of a specified scoring rule; when asked if they knew how the cash component of their bid would be weighted versus the royalty component, bidders indicated they were informed by past auctions, but nonetheless had uncertainty regarding exactly how the decision-makers would choose a winner. The bidders' perspective of the state's choice is better represented by choice probabilities than by a choice rule.

We turn to the auction data to learn more about these choice probabilities. We first look at the illuminating case of two-bidder auctions. Table 1 provides a summary of the patterns observed in bidding and winning. Defining a dominant bid as one which has both a higher cash payment and a higher royalty than the competing bid, we see that there is a dominant bid 62% of the time. Moreover, when there is a dominant bid, the state selects it as the winner 99% of the time. When there is no dominant bid, the bid with higher cash wins 60% of the time, and the bid with higher royalty wins 40% of the time, indicating that the state's choice is not lexicographic.

It is interesting that there is a dominant bid the majority of the time. This pattern is inconsistent with models in which a firm bids a higher royalty because it is relatively cash-poor or because it is more risk-averse. In Louisiana, the bidders bidding higher royalties are simultaneously bidding higher cash, on average.

It is also intriguing that the state strongly favors dominant bids. After all, a high royalty is not an unambiguous good, as it depresses the firm's incentive to develop the lease. The cash-royalty auction here is clearly implemented in a manner that is different from the auctions proposed by McAfee and McMillan (1987) and Laffont and Tirole (1987). Those auctions involve a menu of contracts where cash and royalty are negatively correlated, and the principal would award the contract to the lowest-royalty/highest cash bid. As McAfee and McMillan (1987) and Laffont and Tirole (1987) are concerned with one-dimensional bidder types only, a departure from their optimal mechanism need not indicate suboptimality here. Nonetheless, this gives us some reason to think Louisiana's mechanism may not be optimal.

Figure 2 visualizes the state's choice probabilities using data from all auctions with two or more bids. Regardless of the number of bids, the state's choice of winner among them contains information about pairwise choices that the state made. For instance, if bid A is chosen over bids B and C, we learn about two pairwise choices: that A is chosen over B and that A is chosen over C. Figure 2 plots each of these pairwise choices, with each pair generating two points to plot: the winner (blue circle) and the loser (red triangle). The x and y coordinates of the point represent the royalty and cash components of the bid in relation to the competing bid, so points in the right quadrants have higher royalty than the competing bid, points in the upper quadrants have higher cash, and points in the upper-right quadrant have both higher royalty and higher cash. The transition from red triangles in the lower-left to blue circles in the upper-right visualizes the increased probability of winning as a bid moves in that direction. Moreover, starting at any point in the plot and moving up or to the right, the transition to more winning strongly suggests that the probability of being chosen is increasing in both cash and royalty. The plot also confirms that, even within auction, cash and royalty components are positively correlated; more often than not, one competing bid dominates the other on both dimensions.

2.4 Summary

The data suggest some important features for which a model of these auctions should account. First, the auctioned contract is a real option; the winning firm has a choice

Figure 2: State's choice patterns

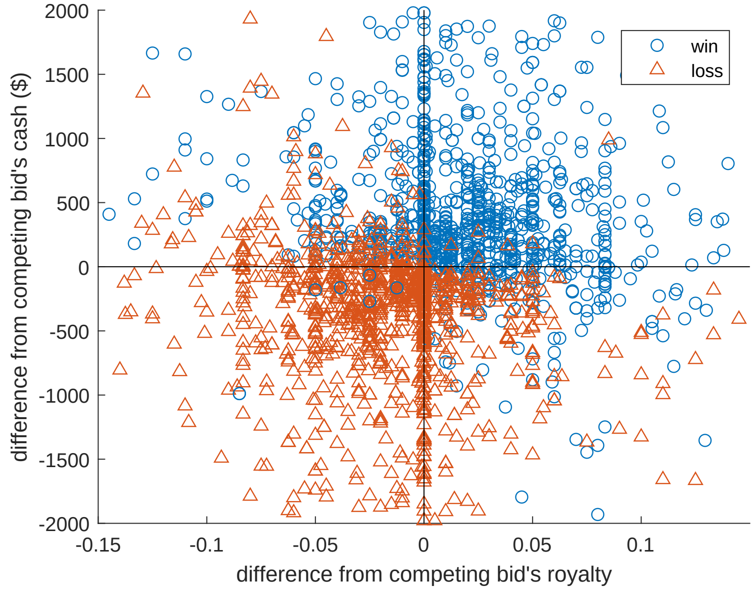

of whether or not to develop it. Moreover, one of the bid components (the royalty) affects the firm's incentives to exercise the option, so the probability of option exercise and the value of the contract (or auction item) to the firm are endogenously affected by the royalty he bids. Second, the bids are two-dimensional, and the distribution of bids in Figure 2 is strongly suggestive of bidders having multidimensional rather than one-dimensional types. Third, there is no scoring rule, so we cannot rely on properties of specific scoring functions. However, much can be learned - by bidders and by the econometrician - from observations of the state's choice over many auctions. In particular, we see that the probability of winning the auction increases in both the cash and royalty components of the bid, and we should not rely on the current mechanism being an optimal one.

3 Model

In this section, we build a model of multidimensional auctions of contracts that incorporates the important features observed in the data. For clarity of exposition, we abstract away from heterogeneity across auction items here, and defer their treatment to the section on estimation.

3.1 Setup

Auction The principal seeks to sell a contract to one of n firms competing to win it. The contract grants the contracted firm the right, but not the obligation, to engage in revenue-generating production of a commodity. Each competing firm makes a two-dimensional bid, (a, b), which leads to a probability of winning given by $P(a, b)$. This probability $P(\cdot, \cdot)$ is increasing in both arguments. The bid component a represents a royalty rate on the firm's revenue to be paid to the principal, and b represents a cash payment. The firm must pay b upon winning the contract regardless of its post-auction actions, but it pays the royalty rate a only if it generates revenue; i.e. the royalty is a contingent payment.

Bidders The bidding firms have two-dimensional types $\vec{\theta} = (\theta_1, \theta_2)$. The type component θ_1 represents the bidder's expected production volume. Since θ_1 is an expected value, it allows for uncertainty while incorporating the productivity of the firm, e.g. well and field design, recovery rate, and other differences in production among firms. The type component θ_2 represents the firm's economic cost of exercising the option to produce, including opportunity costs. Upon winning the contract, a firm decides whether or not to exercise the option. If it does exercise the option, it expects to produce θ_1 units of the good at a cost of θ_2.

Value of the contract From a bidder's perspective, the contract is essentially an option; it is an option to obtain $(1-a)\theta_1$ expected units of the good at a cost of θ_2, where $(1-a)$ represents the portion of revenue kept by the firm after paying royalties. Critically, the bidder's value for the contract must internalize the fact that its probability of exercise is less than one and depends on the future price of the good, which is uncertain at the time of bidding. For example, suppose the contract is near expiration and the price of the good is s. Then the firm will only exercise the option if $s(1-a)\theta_1 > \theta_2$; that is, if its share of revenue after paying royalties to the principal exceeds the cost of option exercise. In the language of options, θ_2 plays the role of a "strike price".

We borrow from the option pricing literature to define a closed-form approximation of the value of such an option. In particular, we value the contract by adapting insights from option models such as the Black-Scholes-Merton model and Black (1976), which assume prices follow a geometric Brownian motion with constant and known volatility.

It is well known that closed-form expressions of value can be derived for European options, which may only be exercised at expiration, but not for American options, which may be exercised at any time until expiration. Therefore, as is common when studying properties of options,[6] we approximate option value using the closed-form of a European option to deduce theoretical properties of the two-dimensional bidding game; later, we use numerical valuation procedures developed for American options to check the robustness of our empirical results.

Let t be the duration in years until the option expires, S the spot price of the good at the time of auction, σ the volatility of spot prices, r the continuously compounded one year risk-free interest rate, and $\Phi(\cdot)$ the standard normal cdf. We assume spot prices follow a geometric Brownian motion with volatility σ and zero expected return after adjusting for inflation. Upon production, a firm receives the spot price at the time of production for each unit produced. Then a bidder's value for the option at the time of its auction is given by

$$V(a,\vec{\theta}) = e^{-rt}[\underbrace{S(1-a)\theta_1}_{\substack{\text{firm's share}\\ \text{of E(revenue)}}}\Phi(x) - \theta_2 \underbrace{\Phi(x-\sigma\sqrt{t})}_{\text{Pr(exercise)}}], \tag{1}$$

where

$$x \equiv \frac{\ln(S(1-a)\theta_1/\theta_2) + \sigma^2 t/2}{\sigma\sqrt{t}}. \tag{2}$$

The expression $\Phi(x-\sigma\sqrt{t})$ represents the ex-ante probability of option exercise, or the probability that the price of the good will be high enough to make the firm's share of revenue exceed its cost. This helps us understand the expression for $V(a,\vec{\theta})$. Intuitively, if option exercise occurred exogenously with probability $\Phi(x-\sigma\sqrt{t})$, the value of the option would instead be $e^{-rt}(S(1-a)\theta_1 - \theta_2)\Phi(x-\sigma\sqrt{t})$: the expected profit from exercise times the probability of exercise, with some present-value discounting for time. However, option exercise is in fact endogenous and occurs when the price of the good is higher than some threshold. So the expected price conditional on exercise is actually higher than the unconditional price, and therefore $S(1-a)\theta_1$ is multiplied by $\Phi(x)$ rather than $\Phi(x-\sigma\sqrt{t})$ to account for this. The model is most

[6] Hull (2017) explains that "some of the properties of an American option are frequently deduced from those of its European counterpart."

Table 2: Effect on option value and probability of exercise

Effect of	on $V(a, \vec{\theta})$	on Pr(exercise)
θ_1	+	+
θ_2	-	-
a	-	-
S	+	+
σ	+	ambiguous
t	+	ambiguous

similar to the commodity option model of Black (1976), but unlike Black (1976) it uses spot prices rather than futures prices because exercising the option results in receiving the commodity itself rather than a futures contract.

Adapting well-known results from the option-pricing literature, Table 2 summarizes how the value of the option and the probability of exercising it are affected by the constituent parameters. Holding all else constant, higher productivity type θ_1 has a positive effect on both option value and exercise probability, while a higher cost type θ_2 has a negative effect, in line with intuition. Also as expected, a higher royalty rate a has a negative effect on both option value and exercise probability.

A few aspects of the model deserve further discussion. First, the model focuses on the extensive margin decision of firms, i.e. of whether they exercise the option, which is a one-shot decision. This is motivated by the fact that a majority of leases are not developed, so the extensive margin decision of firms is the one of first-order interest. It is also motivated by the fact that in conventional oil production, fixed costs are much larger than marginal costs.[7]

Second, the model can be consistent with both private and interdependent value paradigms regarding production volume, depending on the interpretation of θ_1. From the literature on one-dimensional interdependent-value auctions, we know that bidders with primitive information s effectively best-respond to the observed bid distribution "as if" their type was $\phi(s)$,[8] where the function $\phi(s)$ preemptively adjusts s in anticipation of the winner's curse. If the winner's curse adjustment in a multidimensional auction is also characterized by a function $\phi(s)$, then interpreting $\theta_1 = s$ is consistent with private values, and interpreting $\theta_1 = \phi(s)$ is consistent with interde-

[7] As Kellogg (2014) explains in section I.A, operating expenses for maintenance and pumping are small relative to drilling costs, and drilling costs are almost completely sunk.

[8] Specifically, $\phi(s) = \mathbb{E}[\text{production}|S_i = s, \max_{j \neq i} S_j = s]$ for a one-dimensional auction.

pendent values. The theory literature so far does not inform us as to the exact form of the winner's curse adjustment $\phi(\cdot)$ in a multidimensional auction, however; this remains an open question in the literature.

3.2 Cash-royalty bidding

We now turn our attention to properties of the two-dimensional bidding problem. Given his type (θ_1, θ_2), a bidder chooses the two components of his bid (a, b) to maximize his expected utility from the auction. This maximization problem can be written as

$$\max_{a,b}[V(a, \vec{\theta}) - b]P(a, b).$$

To gain an intuitive understanding of the maximization problem, it helps to break it into two steps: in step 1, we consider which combination of cash and royalty the bidder would choose in order to achieve an arbitrary probability of winning, p, since there are in general many elements in the set $\{(a, b)|P(a, b) = p\}$. Then given step 1, we consider the bidder's choice of p. Conceptually separating the bidder's choice of p from how it chooses to achieve that p will yield insights on adverse selection.

Step 1: Bidder's choice of (a, b) for arbitrary p

Now consider said step 1. Fixing p, the bidder chooses its preferred bid from the set $\{(a, b)|P(a, b) = p\}$. Since $P(a, b)$ is increasing in both arguments, a choice of a immediately determines the b required to achieve $P(a, b) = p$. Hence, we can define $b(a, p)$ as a function of a and p and focus on the bidder's choice of a. The bidder's choice of royalty a given p is then

$$a(p, \vec{\theta}) \equiv \max_a \quad \underbrace{V(a, \vec{\theta}) - b(a, p).}_{\text{payoff conditional on winning}}$$

Since the probability of winning is fixed at p, the bidder chooses royalty a to maximize his payoff conditional on winning, which is the value of the contract at that royalty minus the cash payment required to achieve p with that royalty. This yields the first-order condition

$$V_1(a, \vec{\theta}) = b_1(a, p). \tag{3}$$

where a subscript 1 denotes a partial derivative with respect to the first argument. Intuitively, the first-order condition balances the marginal cost of a higher royalty against the marginal benefit. On the one hand, bidding a higher royalty allows a bidder to bid less cash (right-hand side of (3)), but it also reduces the value of the contract to him (left-hand side of (3)).

The proof of the following proposition shows that bidding a higher royalty percentage is less costly for undesirable "weak" types - those with low productivity (θ_1) and/or high cost (θ_2) - than for "strong" types. Mathematically, the effect of a higher royalty on contract value, $V_1(a,\vec{\theta})$, becomes more negative with θ_1 and less negative with θ_2. This has important implications for adverse selection; the royalty component of bidding effectively provides weak types a cheaper currency with which to bid. Naturally, this leads strong types to favor a low-royalty high-cash bid and weak types to favor a high-royalty low-cash bid from among the set $\{(a,b)|P(a,b)=p\}$.

Proposition 1. *The marginal effect of royalty a on contract value, $V_1(a,\vec{\theta})$, is decreasing in θ_1 and increasing in θ_2. As a result, the bidder's choice of royalty as a function of p and $\vec{\theta}$, $a(p,\vec{\theta})$, is decreasing in θ_1 and increasing in θ_2.*

Note that the royalty a is a function of both θ_1 and θ_2. The same goes for the cash payment b. This is not a model in which each bid component is purely a function of one type component.

Step 2: Bidder's choice of p given $a(p,\vec{\theta})$

Now consider step 2, the bidder's choice of probability of winning, p. Knowing from step 1 that the bidder will choose to bid royalty $a(p,\vec{\theta})$ and cash $b(a(p,\vec{\theta}),p)$ given p, we define the payoff conditional on winning as a function of p and $\vec{\theta}$ as follows:

$$\pi(p,\vec{\theta}) \equiv V(a(p,\vec{\theta}),\vec{\theta}) - b(a(p,\vec{\theta}),p).$$

Then the bidder chooses p to maximize his expected utility from the auction, $\pi(p,\vec{\theta})p$. Formally, the maximization problem and the corresponding first-order condition are

$$p(\vec{\theta}) \equiv \max_p \pi(p,\vec{\theta})p,$$

$$\frac{1}{p} = -\frac{\pi_1(p,\vec{\theta})}{\pi(p,\vec{\theta})}.$$

As is standard in auctions, the bidder would like to increase his probability of winning, but achieving a higher p requires a more competitive bid, which decreases his payoff conditional on winning. The bidder's choice of p optimally balances these two forces.

As we look ahead to bringing our bidding model to data, it is helpful to rewrite the first-order conditions from steps 1 and 2 in terms of observables. If the observed bid corresponds to the bidder's optimal choice, we can replace $a(p, \vec{\theta})$ with observed a, $b(a(p, \vec{\theta}), p)$ with observed b, and p with $P(a, b)$. Further, we replace $\pi(p, \vec{\theta})$ with its definition, use $b_1(a, p) = -P_1(a, b)/P_2(a, b)$ by properties of implicit derivatives, and use $\pi_1(p, \vec{\theta}) = -b_2(a(p, \vec{\theta}), p) = -1/P_2(a, b)$ by the envelope theorem. Then the bidder's choice of two-dimensional bid (a, b) as a function of his two-dimensional type (θ_1, θ_2) is characterized by the following system of two first-order conditions:

$$V_1(a, \vec{\theta}) = -\frac{P_1(a, b)}{P_2(a, b)}, \tag{4}$$

$$V(a, \vec{\theta}) = b + \frac{P(a, b)}{P_2(a, b)}. \tag{5}$$

In particular, equation (5) is familiar in form; it resembles the first-order condition for first-price auctions as written in Guerre, Perrigne, and Vuong (2000).

In auctions with a known scoring rule, a multidimensional auction is outcome-equivalent to an auction in which bidders bid on the scalar score only and the highest score wins. The existing literature has exploited this idea to characterize equilibrium bidding for a restricted class of scoring auctions, including auctions with quasi-linear scoring rules. To our knowledge, all work that characterizes equilibrium in multidimensional auctions has relied on this scalar reduction. In the absence of a known scoring rule, we do not characterize equilibrium bidding. Rather, we characterize the bidder's optimal bid in response to the probability of winning $P(a, b)$, which is observed by bidders and the econometrician through auction outcomes. This indirect approach is in the spirit of Guerre, Perrigne, and Vuong (2000). It is also related to Larsen and Zhang (2018), who exploit observed outcomes of agents' actions to identify their types when the precise rules of the game are unknown.

4 Identification

We show that the type (θ_1, θ_2) associated with each observed bid (a, b) is identified nonparametrically from the observables: the joint distribution of bids $G(a, b)$ and whether each bid wins or loses.

Identification requires that there be a unique (θ_1, θ_2) that rationalizes bid (a, b) given the observed patterns of winning and losing. Now consider equations (4) and (5) that characterize the mapping between types and optimal bids. In these equations, $V(\cdot, \cdot)$ is a known function, and $P(\cdot, \cdot)$, the probability of winning as a function of the bid, is identified immediately from observations of bids and whether they win. Then θ_1 and θ_2 are the only remaining unknowns. Equations (4) and (5) thus constitute a system of two nonlinear equations in two unknowns, and we can show identification by proving that the system has a unique solution. The proof of the following proposition does exactly this.

Proposition 2. *The two-dimensional bidder type* (θ_1, θ_2) *associated with bid* (a, b) *is identified nonparametrically.*

Intuitively, the separate identification of θ_1 and θ_2 implies that bidders' types cannot be reduced to scalar representations in this model. To provide a contrasting example, suppose royalties were levied on profit instead of revenue and development was mandatory. In this case, $V(a, \vec{\theta}) = (1 - a)(S\theta_1 - \theta_2)$ and $V_1(a, \vec{\theta}) = -(S\theta_1 - \theta_2)$. Then in (4) and (5), $\vec{\theta}$ could be replaced by the scalar value $(S\theta_1 - \theta_2)$ without loss, and all (θ_1, θ_2) pairs in the set $\{(\theta_1, \theta_2)|(S\theta_1 - \theta_2) = v\}$ would be observationally equivalent; i.e. separate identification of θ_1 and θ_2 would fail. In reality, royalties are levied on revenue not profit, so the distinctive roles played by θ_1 versus θ_2 in determining the bid cannot be reduced to or represented by the bidder's overall value for the contract. Continuing with the mandatory development example for sake of intuition, royalty being levied on revenue implies $V(a, \vec{\theta}) = (1 - a)S\theta_1 - \theta_2$ and $V_1(a, \vec{\theta}) = -S\theta_1$. Then, the separate identification of θ_1 and θ_2 is apparent given the system of (4) and (5). The appendix provides a proof for the full model with optional development, which shares some of this intuition but is mathematically more subtle.

5 Estimation

We develop a multi-step nonparametric estimation procedure, which entails estimating the probability of winning $P(a, b)$ and subsequently solving for (θ_1, θ_2) using (4) and (5).

Given a large data sample, $P(a, b)$ can be estimated directly via a nonparametric regression of win/loss dummies on bid components (a, b). At the same time, the probability of winning is actually a structural composite of two components: the bid distribution and the state's pairwise choice probability, described earlier in section 2.3. We exploit this structure to estimate $P(a, b)$ well in a finite sample. This approach has the additional advantage of being more explicit about what generates $P(\cdot, \cdot)$ so that its origins are not a blackbox. Before we get to the details of this estimation, we first address heterogeneity across auction items.

5.1 Auction-level heterogeneity

All empirical auction studies of non-homogeneous goods must address auction-level heterogeneity before estimating the bid distribution a bidder competes against. No two auction items are exactly the same, so the heterogeneity needs to be controlled for in some way. Let z represent the vector of descriptive auction-level covariates that differ among auction items. As nonparametric conditioning on z would suffer from the curse of dimensionality, we propose the following.

For each auction item, we predict the mean of a conditional on z and the mean of $\ln b$ conditional on z using "leave-one-out" regression of bid components a and $\ln b$ on z, respectively. Leave-one-out here refers to not using own-auction bids in the regression, rather using a, b, z data from all the other auctions. Then, for purposes of estimating the bid distribution only, we convert each a to its deviation from the (auction-specific) conditional mean royalty, and convert each b to its log deviation from the conditional mean cash bid. In other words, the bid's competitive position is characterized by how many more percentage points it bids in royalty and how much more (in logs) it bids in cash relative to the conditional mean. This procedure normalizes bids with respect to the estimated conditional mean and allows pooling of bids across heterogeneous auctions when estimating the bid distribution $G(a, b)$.

5.2 Estimation of the probability of winning $P(a, b)$

To estimate $P(a, b)$, we first estimate the bivariate bid distribution $G(a, b)$ via kernels. Second, we estimate the state's pairwise choice probability $C(a, b, a', b')$ via sieve maximum likelihood, where $C(a, b, a', b')$ is defined as the probability that the state chooses bid (a, b) over a competing bid (a', b'). Finally, we estimate $P(a, b)$ as the probability that bid (a, b) is chosen over n competing bids drawn from $G(a, b)$.

First, let $(\tilde{a}, \tilde{b})$ denote the bids normalized according to section 5.1. We estimate the bivariate distribution of $(\tilde{a}, \tilde{b})$, denoted $\tilde{G}(\cdot, \cdot)$, nonparametrically via kernel estimation. For bandwidths, we use Silverman's rule of thumb, which for two dimensions is h_i = (sample size)$^{-1/6}$(standard deviation of ith variable). Then, if $\bar{a}(z)$ and $\bar{b}(z)$ are the mean bid components conditional on z as described in section 5.1, we can define $G(\cdot, \cdot; z)$ as a simple transformation of $\tilde{G}(\cdot, \cdot)$: $G(a, b; z) \equiv \tilde{G}(a - \bar{a}(z), \ln b - \ln \bar{b}(z))$. Thus we estimate $G(a, b)$ conditional on z. As is usual with auctions, the bid distribution should be estimated fixing the number of bidders n, as it is likely to change with n.

Second, we estimate the state's pairwise choice probability $C(a, b, a', b')$ nonparametrically via sieve maximum likelihood. Recall from section 2.3 that "pairwise" does not refer to the number of bidders n equaling two; each n-bidder auction yields $n - 1$ observed pairwise choices made by the state between the winning bid and each of the losing bids, and all of these observations can be used for estimation. For each observed pairwise choice between (a, b) and (a', b'), the components of the pair of bids are the inputs to $C(\cdot, \cdot, \cdot, \cdot)$, and an indicator for (a, b)'s win, w, is the binary outcome. We reduce the dimensionality of this function by assuming the state compares two bids based on their component-wise differences; that is, $C(a, b, a', b') = C(a - a', \ln b - \ln b')$. To condition this function on auction-level heterogeneity, we include an index $q(z)$ of item quality as an additional argument and estimate $C(a - a', \ln b - \ln b', q(z))$. The quality index $q(z)$ is formed as the predicted value of a regression of $\ln(b)$ on auction covariates z.

The log-likelihood, over all observed pairwise choices, of observed outcome w given bids (a, b) and (a', b') is

$$\sum_{k=1}^{K} w_k \ln C(a_k - a'_k, \ln b_k - \ln b'_k) + (1 - w_k) \ln[1 - C(a_k - a'_k, \ln b_k - \ln b'_k)],$$

where K is the number of observed pairwise choices. Sieve maximum likelihood approximates $C(\cdot,\cdot)$ with sieves - we use multivariate Bernstein polynomials - and estimates the polynomial parameters that maximize the likelihood. For regularity, we restrict $C(\cdot,\cdot)$ to increase in its first two arguments, reflecting the state's preference for higher bid components as observed in section 2.3. Since the domain of Bernstein polynomials is $[0,1]$, we convert each argument of $C(\cdot,\cdot)$ to its quantile in the observed distribution of that argument for purposes of computing the polynomials.

Finally, having estimated $G(\cdot,\cdot)$ and $C(\cdot,\cdot)$, we estimate $P(\cdot,\cdot)$ as their structural composite. Using the case of two-bidder auctions as an illustrative example, the probability of winning with bid (a,b) is the expectation of pairwise choice probability over the distribution of competing bids,

$$P(a,b) = \int C(a-a', \ln b - \ln b')dG(a',b').$$

Subsequently, the derivatives of $P(a,b)$, including $P_1(a,b)$ and $P_2(a,b)$, are computed numerically.

5.3 Estimation of (θ_1, θ_2)

The proof of Proposition 2 provides a constructive and detailed description of how to solve the system of equations (4) and (5) mathematically. This section provides a few additional details regarding practical implementation.

As explained in Guerre, Perrigne, and Vuong (2000) regarding bid inversions that exploit density estimates, some trimming is necessary to correct for the boundary effects of density estimation. Observations are relatively sparse in the tails of the bid distribution, so we follow Li, Perrigne, and Vuong (2000) in our trimming; we estimate (θ_1,θ_2) for values of $(\tilde{a},\tilde{b})$ that are at least two bandwidths away from the extrema, i.e. $(\tilde{a},\tilde{b}) \in [\tilde{a}_{min} + 2h_{\tilde{a}}, \tilde{a}_{max} - 2h_{\tilde{a}}] \times [\tilde{b}_{min} + 2h_{\tilde{b}}, \tilde{b}_{max} - 2h_{\tilde{b}}]$.

Meanwhile, the expression for $V(a,\vec{\theta})$ in (1) shows that solving the system (4) and (5) for (θ_1,θ_2) requires inverting the standard normal cdf $\Phi(\cdot)$. While $\Phi(\cdot)$ is a strictly increasing function, it is nearly flat when the cdf value is close to 0 or 1. Therefore, we estimate (θ_1,θ_2) if $\Phi(x-\sigma\sqrt{t}) \in (\epsilon, 1-\epsilon)$, where $\epsilon = 10^{-4}$.

6 Application to Louisiana cash-royalty auctions

6.1 Details

We bring our model to the Louisiana cash-royalty auctions and apply our estimation method. Our sample, described in section 2, constitutes 977 auctions with at least 2 bidders. We use the pairwise choices observed from all of these auctions to estimate the state's pairwise choice probability $C(\cdot,\cdot)$. Meanwhile, the bid distribution $G(\cdot,\cdot)$ should be estimated for a fixed number of bidders. 772 out of 977 auctions in our sample have two bidders, so we estimate $G(\cdot,\cdot)$ and $P(\cdot,\cdot)$ conditional on 2 bidders using these 772 auctions. After estimating $P(\cdot,\cdot)$, solving the system of first-order conditions in (4) and (5) requires historical data on implied volatilities, which are available from 1987. Therefore, the final step of solving for (θ_1, θ_2) given $\hat{P}(\cdot,\cdot)$ is completed for bids from 1987-2003, a period encompassing 250 auctions with $n = 2$ in our sample.

To obtain the one year risk-free interest rate r, we take nominal 1-year treasury rates provided by FRED and convert them to real rates via application of the Fisher equation, using percentage changes in the GDP implicit price deflator as the measure of inflation. All variables subject to inflation are inflation-adjusted using the GDP implicit price deflator. For S, the spot price of oil, we use the West Texas Intermediate (WTI) spot price provided by FRED. For σ, the volatility of S, we use the implied volatility of 1-month crude oil futures, i.e. the expected volatility implied by contemporary crude oil option prices. For that purpose, we purchased historical prices of crude oil options from the CME Group; further details on deriving implied volatility are provided in the appendix.

To control for auction-level heterogeneity as described in section 5.1, our auction covariates z include the spot price S, the interest rate r, the log of tract acreage and its square, an indicator for royalty recipient type,[9] auction year fixed effects, and two variables to account for geological and geographic heterogeneity, constructed as follows. First, we use historical production data from Drillinginfo to compute a production index for each township in Louisiana; it is the log of barrel-of-oil-equivalents (BOE) produced per well in each township. Second, we perform a local quadratic regression of observed bid components on the geographic coordinates of the associated

[9]Specifically, tract kind code 4 indicates that the royalty recipient is a state agency other than the Department of Natural Resources.

Figure 3: Estimated bivariate bid density $\hat{\tilde{g}}(\tilde{a}, \tilde{b})$

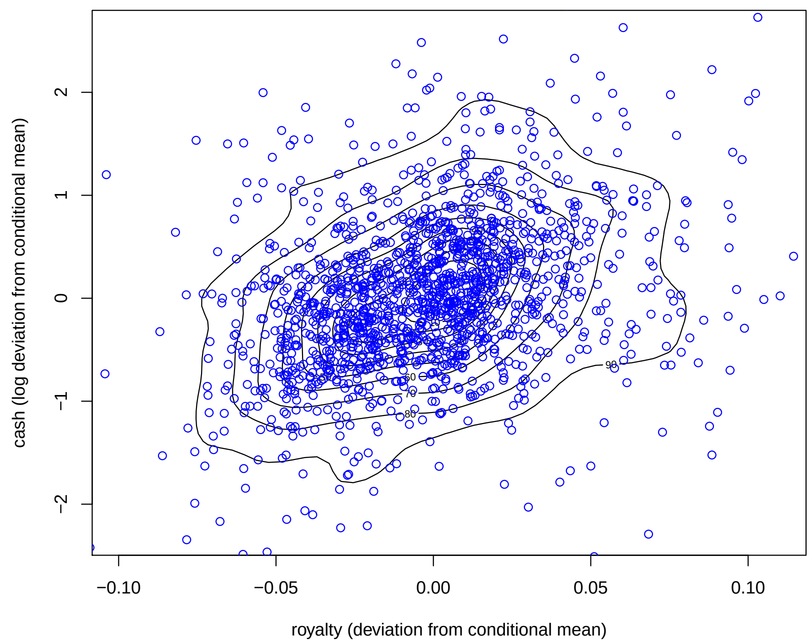

lease's township; this surface fit is done once for each auction, excluding own-auction bids, to produce a fitted royalty index and cash payment index for each auction based on geographic location. These indices are included as auction covariates to capture geological and geographic heterogeneity.

6.2 Estimation results

Figures 3 and 4 plot level contours of the estimated bid density and pairwise choice probability, respectively, along with the observations used to estimate them. Figure 3 shows that the bid distribution is densest near the conditional mean bid and gets sparser farther from the mean. Most bids are within 5 percentage points of the conditional mean royalty and within a log deviation of 1 from the conditional mean cash payment. The contours in Figure 4 show the transition from low probability to high probability of being favored by the state as bid components increase, which directly follows the transition from losses (red triangles) to wins (blue circles) in the underlying data.

Figure 4: Estimated pairwise choice probability $\hat{C}(a - a', \ln b - \ln b', q(z))$ at median $q(z)$

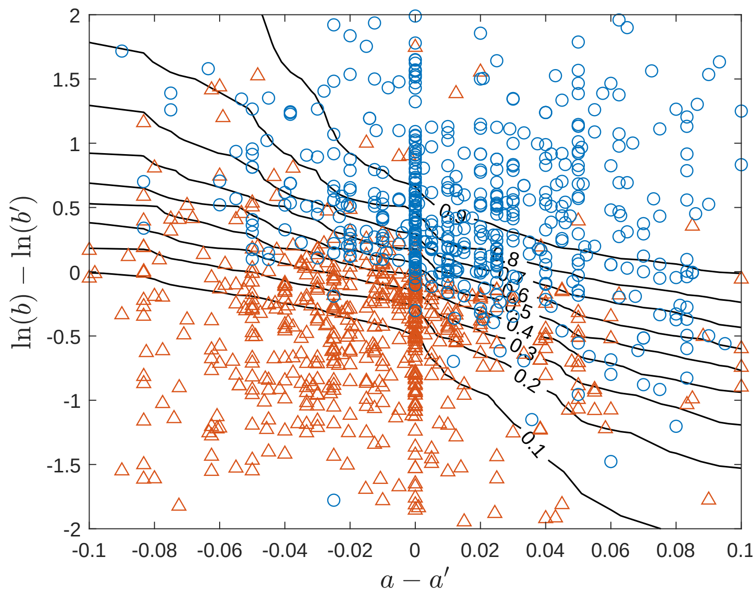

Figures 5 and 6 plot the densities of estimated productivity type $\hat{\theta}_1$ and per-unit cost $\hat{\theta}_2/\hat{\theta}_1$, respectively. Note that these are raw values across all auctions in the estimation sample; they are not conditioned on any particular value of auction covariates. The ratio $\hat{\theta}_2/\hat{\theta}_1$ represents the bidder's cost per unit of production (barrel), which is a more intuitive measure of cost than $\hat{\theta}_2$ on its own. Figure 6 includes for comparison a red dotted line representing the average of $(1 - a)\times$oil price in the estimation sample; $(1 - a)\times$oil price is the revenue a firm would keep per barrel if it developed the tract. Supposing $(1 - a)\times$oil price were fixed, firms with per-unit cost to the left of the dotted line would exercise their option, while firms to the right of it would not. As there is substantial mass to the right of the red line, we begin to see why so many leases are not developed.

According to the model used in (1), the ex-ante probability of option exercise is equal to $\Phi(x - \sigma\sqrt{t}) = \Phi([\ln(S(1 - a)\theta_1/\theta_2) - \sigma^2 t/2]/(\sigma\sqrt{t}))$. "Ex-ante" refers to the fact that this is the probability before the future price path for the lease is realized. While the comparison of ex-ante to ex-post is not apples-to-apples, we compare the ex-ante predicted probability to the ex-post observed probability for a sense of model fit. For this purpose, we define a lease as "developed" if the Louisiana DNR records receipt of royalties from the lease or records a well attached to the lease. Since

Figure 5: Kernel density of estimated $\ln \hat{\theta}_1$

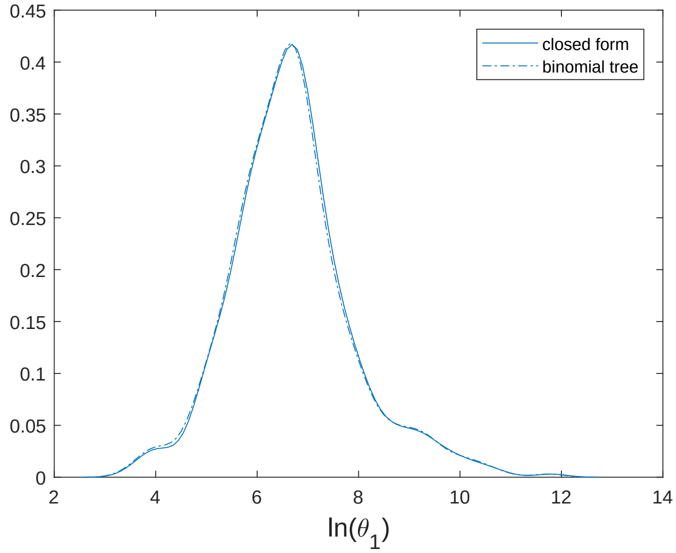

only winning firms decide whether to exercise the option, only winning bids should be used in calculating the ex-ante probability. Based on winning bids submitted for leases where the royalty recipient is the DNR,[10] the average predicted ex-ante probability of exercise is 0.4. For the same set of leases, the average observed ex-post probability of exercise is 0.35. Our estimates of (θ_1, θ_2) are consistent with the low rate of development observed, though we do not use post-auction data in estimation.

6.3 Robustness: valuing leases using a binomial tree

To ensure that the conclusions we draw from our analysis are not sensitive to the European-option approximation underlying the closed-form expression in (1), we also use a numerical procedure known as a binomial tree to value the lease as an American option and re-estimate (θ_1, θ_2) accordingly. A binomial tree begins with a single node, and each node connects to two nodes in the next period; one represents the probability that oil prices will go up, the other the probability that oil prices will go down. At each node, the agent has a choice of whether to exercise the option given the node-specific oil price s. The agent exercises the option at that node if $s(1-a)\theta_1 - \theta_2$ exceeds the continuation value. The value of the node equals $s(1-a)\theta_1 - \theta_2$ if the agent would choose to exercise, and it equals the continuation value otherwise. The

[10]The DNR does not record receipt of royalties if the royalty recipient is a different state agency, indicated by tract kind code 4.

Figure 6: Kernel density of estimated $\hat{\theta}_2/\hat{\theta}_1$

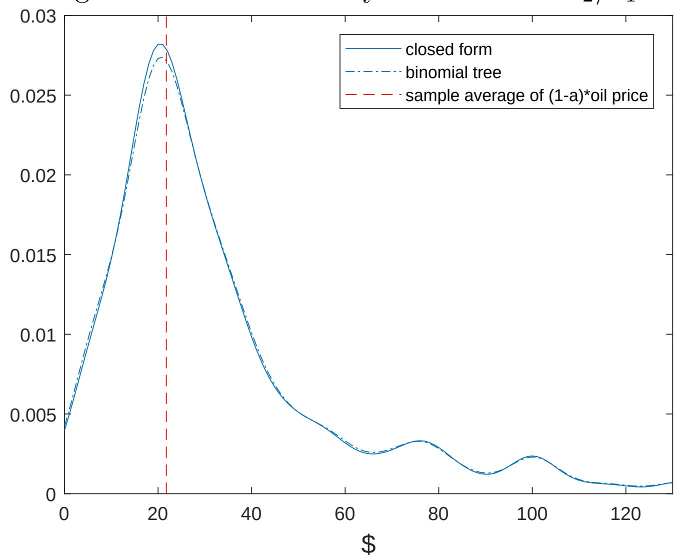

value of the American option is then computed as the present-discounted expected value of all nodes that extend from the first node. Reference books such as Hull (2017) provide details on constructing binomial trees given the same inputs that go into (1).[11] We use a binomial tree with 1000 steps (periods) to value $V(a, \vec{\theta})$ and then use constrained optimization to estimate the (θ_1, θ_2) that best satisfy the system in (4) and (5). The density of $(\hat{\theta}_1, \hat{\theta}_2)$ thus estimated are presented alongside those of the closed-form estimates in Figures 5 and 6. In addition, all counterfactuals in the next section are repeated with the binomial tree valuation to ensure robustness.

7 Counterfactual analysis

7.1 Cash-royalty auctions versus fixed-royalty auctions

Fixed-royalty auctions – the prevailing auction design for public mineral leasing in the U.S. – are a simpler alternative to cash-royalty auctions. In that auction format, the principal fixes a common royalty rate, and bidders bid only on the cash payment amount. Cash-royalty auctions have the potential to yield better outcomes for the principal than fixed-royalty auctions, but how they perform in practice is an empirical question that depends on the details of implementation. Using our estimated

[11] We use code that is adapted from Paolo Zagaglia's MATLAB option pricing package.

structural model, we compare Louisiana's cash-royalty auction to the fixed-royalty benchmark. The model is sufficiently rich to allow comparison in the particulars of allocative distortions, moral hazard, and information rents, the combined effects of which lead to an overall comparison of social surplus and government revenue. To conduct this counterfactual comparison, we simulate fixed-royalty auctions with fixed royalty rates ranging from 0% to 50%. All evaluations are ex-ante with respect to oil price realizations, so results are not tied to a specific path of oil prices.

In the graphs that follow, the solid blue curve represents the fixed-royalty auction. The movement of the solid curve from left to right shows how the quantity of interest is affected by the level of the fixed royalty, which is indicated on the x-axis. Meanwhile, the dotted horizontal line marks the value of the quantity of interest in Louisiana's cash-royalty auction. Since the royalty is not fixed in a cash-royalty auction, the x-axis is irrelevant for the dotted horizontal line. We are interested in comparing the level of the dotted line to the solid curve.

Allocative effects of leasing versus selling In auctions with one-dimensional value types and bids, the allocative effects of royalty leasing relative to an outright sale cannot be meaningfully analyzed. Valuations would be monotonic in one-dimensional bidder types with or without a royalty, so by construction the highest-value bidder under leasing would always coincide with that of an outright sale. Here, generalizing the type space to two dimensions allows us to meaningfully analyze the allocative effects of leasing design. To do this, we first take note of which bidder would have won in an auction where the mineral rights were sold outright, i.e. a zero royalty auction. Using this as the benchmark, we then compute under different designs the fraction of auctions that retain the benchmark allocation.

The blue curve in Figure 7 shows that allocative deviations from the benchmark increase with royalty for fixed-royalty auctions. Intuitively, as royalties are levied on revenue alone, high royalties are especially punishing for bidders who have both high productivity and high cost, compared to bidders who have low productivity and low cost. This leads to flips in the winner's identity as the fixed royalty increases. Meanwhile, we see from the orange dotted line that Louisiana's auction system results in allocative deviations that are more pronounced than any of the fixed-royalty auctions plotted. This is because the endogenous choice of royalty by bidders in Louisiana exacerbates adverse selection. As shown in Proposition 1, royalties are less costly for

Figure 7: Fraction of auctions that achieve socially optimal allocation

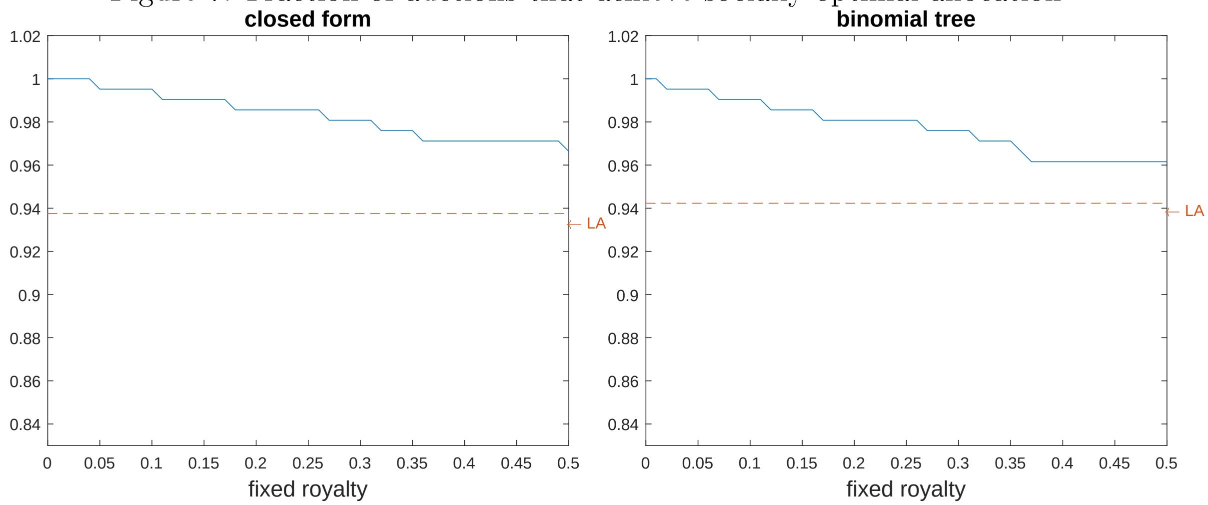

low productivity and/or high cost types and serve as a cheaper currency with which to bid. As a result, these undesirable types win more often than when royalty is fixed by the government.

Moral hazard For the designers of the Louisiana leasing system, one intended benefit of allowing the "market" to determine royalties would have been to avoid the incidence of inflexible fixed royalties on tracts that are not productive enough to support them. Such flexibility could potentially increase the rate of option exercise compared to a rigid fixed-royalty system in which one royalty rate – such as the federal BLM's 12.5% – is applied universally. It would also lift the government's burden of having to determine the fixed royalty. However, we see in Figure 8 that the Louisiana auction does not achieve any special benefit for the probability of option exercise; the average royalty in the Louisiana leasing system is 23%, and the probability of option exercise would have been slightly higher if the government had simply fixed the royalty at this average. As expected, the blue curve shows that the probability of option exercise is decreasing in fixed royalty, as the firm's incentive to exercise the option declines with the portion of revenue transferred to the government.

Social surplus Social surplus in this context is defined as all production revenues minus costs. It incorporates the allocative effects and moral hazard which were discussed above. The designers of Louisiana's leasing system may have intended its flexibility to yield benefits for social surplus. However, as was the case in our investi-

Figure 8: Probability of option exercise

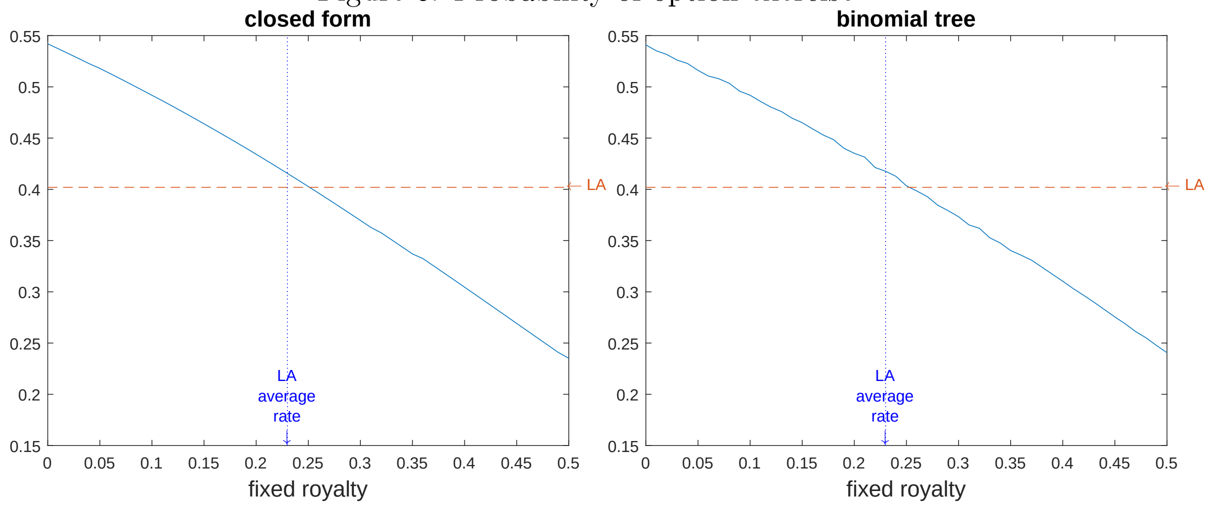

gation of moral hazard, there does not seem to be any special benefit; Figure 9 shows that the social surplus in Louisiana (orange dotted line), which has average royalty of 23%, is slightly lower than what it would have been had the government just fixed the royalty at that rate. This finding contrasts with studies of multi-attribute auctions without moral hazard, such as Lewis and Bajari (2011), in which multi-attribute bidding yields welfare gains over fixing the non-cash component. Meanwhile, as expected given the allocative effects and moral hazard associated with royalties, the blue curve shows that social surplus is declining in royalty.

Cash payment from winning firm to government We now turn our attention to the transfers between agent and principal. Figure 10 plots the cash component of the winner's bid, which is paid by the winning firm to the government. Naturally, the blue curve shows that among fixed-royalty auctions, a higher fixed royalty decreases the value of the lease to the bidder, resulting in lower cash bids. In Louisiana's cash-royalty auction, though, the cash payment is even lower than the average royalty level (23%) would lead us to expect. Specifically, Figure 10 shows that the average cash payment in Louisiana (orange dotted line) is about \$1000 per acre. In a 23% fixed-royalty auction, it would have been \$400 more at about \$1,400 per acre. Compared to having one-dimensional competition in cash only, Louisiana seems to disproportionately weaken cash competition relative to the royalty level it achieves.

Figure 9: Social surplus

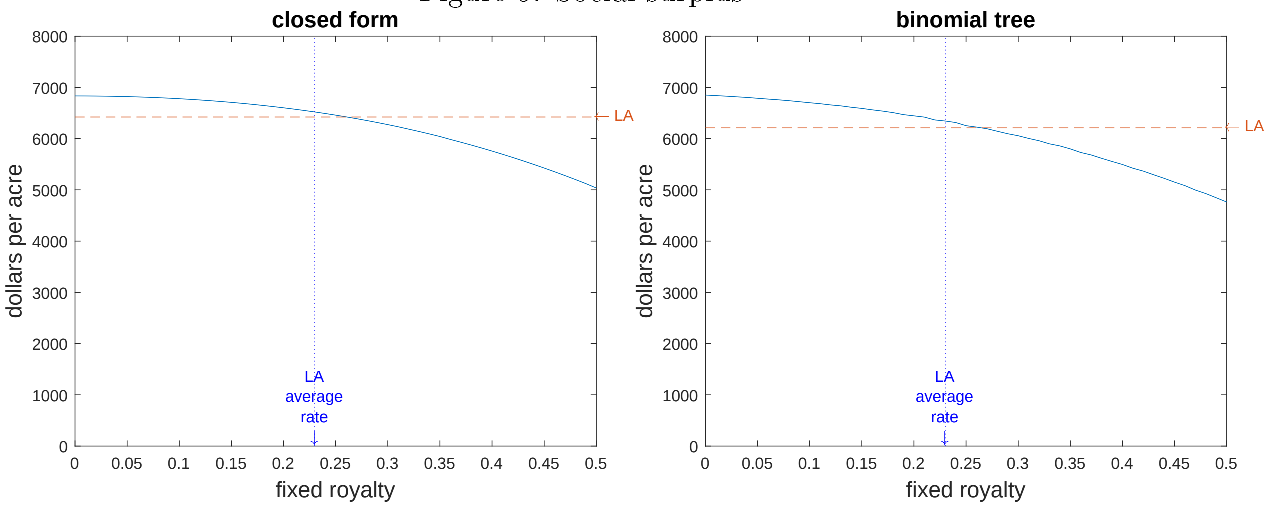

Figure 10: Cash payment from winning firm to government

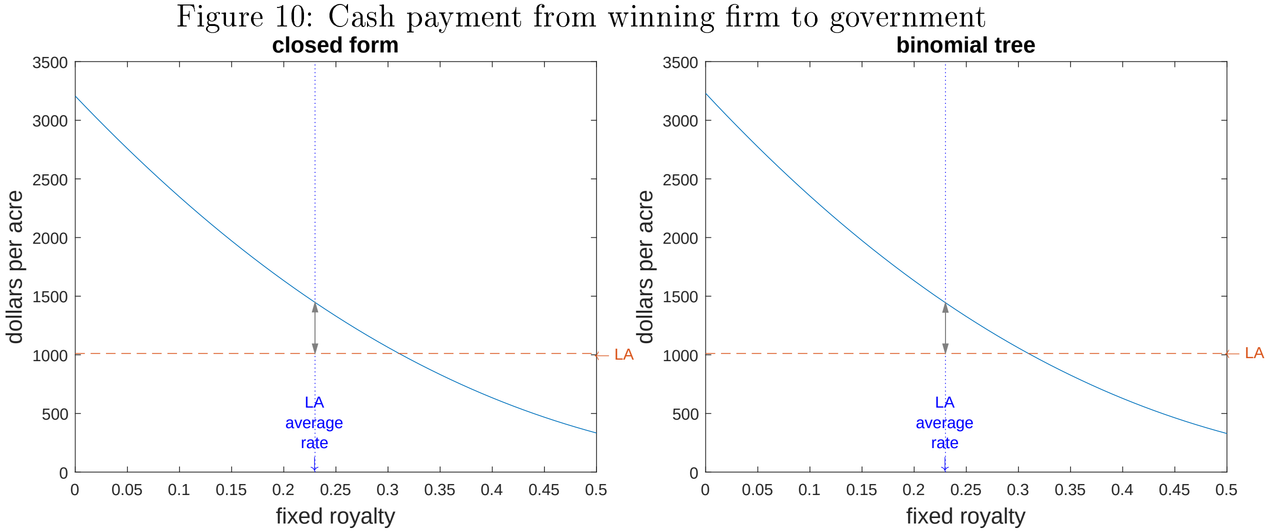

Figure 11: Information rents

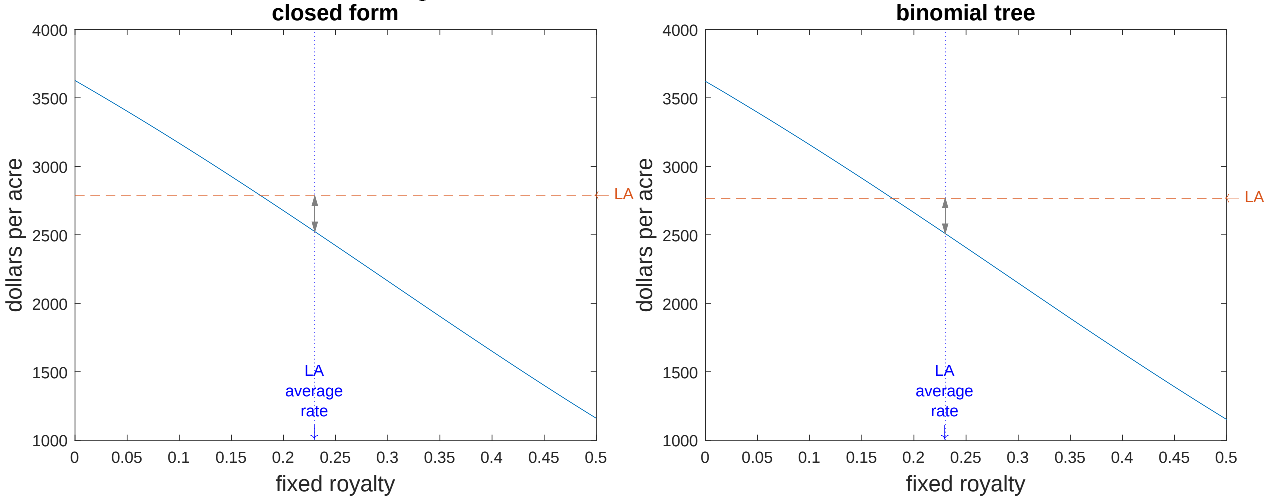

Information rents (winning firm's surplus) To get a sense of the firms' surplus under each auction design, we consider their information rents. Information rents are defined as the firm's ex-ante value for the lease, evaluated at the relevant royalty rate, minus the cash component of its bid. We know from theory that agents' information rents are decreasing in royalty; the more the principal's payoff is tied to the actual outcome, the less room there is for the agent to reap information rents. The blue curve in Figure 11 confirms this. Meanwhile, we see that as a consequence of weaker cash competition as discussed above, bidders in Louisiana reap higher information rents than the average royalty level (23%) would lead us to expect. The winning firm's average surplus is about $2,800 per acre in Louisiana's cash-royalty auction. In a 23% fixed-royalty auction, it would have been $300 less at about $2,500 per acre.

Total government revenue (sum of cash and royalties) Combining allocative effects, moral hazard, and firms' information rents, we now assess the effect of auction design on overall government revenue, defined as the sum of cash payments and ex-ante expected royalties. In line with what we have seen in cash payments and information rents, Louisiana's cash-royalty auction does not perform well in terms of total government revenue, as shown in Figure 12. In fact, a fixed-royalty auction with royalty fixed at any rate between 10% and 50% would yield higher total revenue for the government. This interval includes the federal BLM's fixed rate of 12.5% as well as the 25% rate which is common on privately owned lands. In particular,

Figure 12: Total government revenue

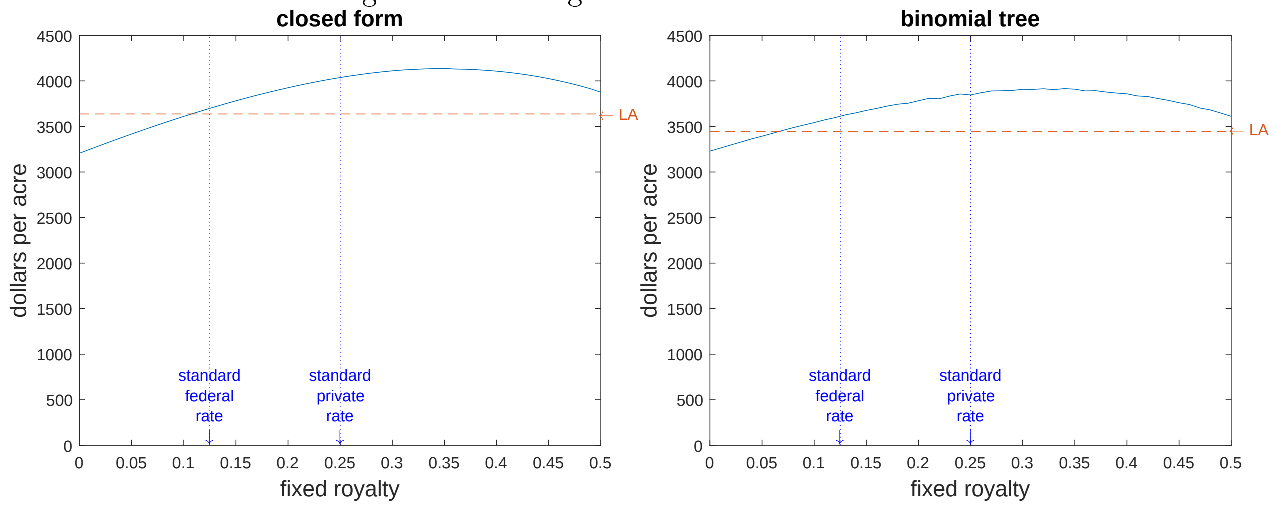

implementing a 25% fixed-royalty auction rather than the current design would yield an increase in revenue from about \$3,600 to \$4,000 per acre.

7.2 Comparison to scoring auctions

Louisiana implements its cash-royalty auction without specifying a scoring rule. If it retained two-dimensional cash-royalty bidding but announced a scoring rule, how would the outcome compare to the current design and to the fixed-royalty auction simulated in the previous section? To address this question, we simulate auctions with quasi-linear scoring rules of the form $s(a, b) = b + \omega S(-1/a^{\rho})$, where ω is the weight the score places on the royalty component, S is the spot price of oil at the time of auction, and ρ represents the curvature of the score with respect to the royalty component. We simulate scoring rules with a number of different curvatures by varying ρ. For each curvature ρ that we simulate, we perform a grid search for the weight ω that would maximize government revenue. The result we present for each ρ is associated with the respective optimal ω.

Figure 13 displays bidding patterns and revenue from the scoring auctions. The x-axis indicates the curvature ρ of the scoring rule, which varies from 1 to 10. The blue dashed-line curve, read by the left y-axis, shows the variance of royalties that bidders choose to bid under each scoring rule. The solid orange curve, read by the right y-axis, shows the total government revenue obtained by each scoring rule. The figure shows that as the curvature ρ of the scoring rule increases, the variance of royalties bid by bidders decreases. As the variance of royalties decreases, total government revenue increases and approaches that of the fixed-royalty auction with optimally chosen royalty. Meanwhile, the simulated scoring auctions all achieve higher revenue than the current cash-royalty auction in Louisiana.

For scoring auctions with no moral hazard, such as price-quality auctions, Asker and Cantillon (2008) prove in their Theorem 6 that the principal is "better off using a scoring auction with a scoring rule that corresponds to his true taste than imposing [the quality level] and selecting the winner on the basis of price only." On the other hand, our simulations suggests that when a bid component – the royalty in our data – induces moral hazard, scoring auctions do worse than imposing a level on the dimension that induces moral hazard. As explained in Proposition 1, whenever royalty can substitute for cash in the principal's selection rule, less desirable types will exploit this tradeoff to the principal's detriment, since promising a high royalty is cheap for these types. Figure 13 implies that the more room there is for such a tradeoff, the worse the outcome is for the principal.

7.3 Effects of lease duration and timing

Having estimated bidders' types and equipped with our option model, we are also in a position to address policy questions beyond auction design. One such question has to do with the duration of the lease. To determine the effects of a longer lease on government revenue and probability of option exercise, we counterfactually simulate fixed-royalty auctions of a 6-year lease, doubling the observed duration of 3 years. Recall from Table 2 that the effect of contract duration on option value is positive, while its effect on option exercise is theoretically ambiguous. The top plot in Figure 14 shows that for Louisiana, doubling the lease duration would in fact decrease the ex-ante probability of option exercise at the most common royalty values. The longer duration does still increase option values and hence cash bids, leading to the bottom plot in Figure 14 which shows an overall increase in government revenue. However, the revenue increase is less than proportional to the duration increase.

Meanwhile, better exploiting fluctuations in oil prices is a policy that may yield substantial benefits. As Louisiana has control over auction offerings, they can withhold leases when oil prices are low and release them when oil prices are higher. Figure 15 counterfactually simulates what government revenue would have been had oil prices

Figure 13: Counterfactual simulation of scoring auctions

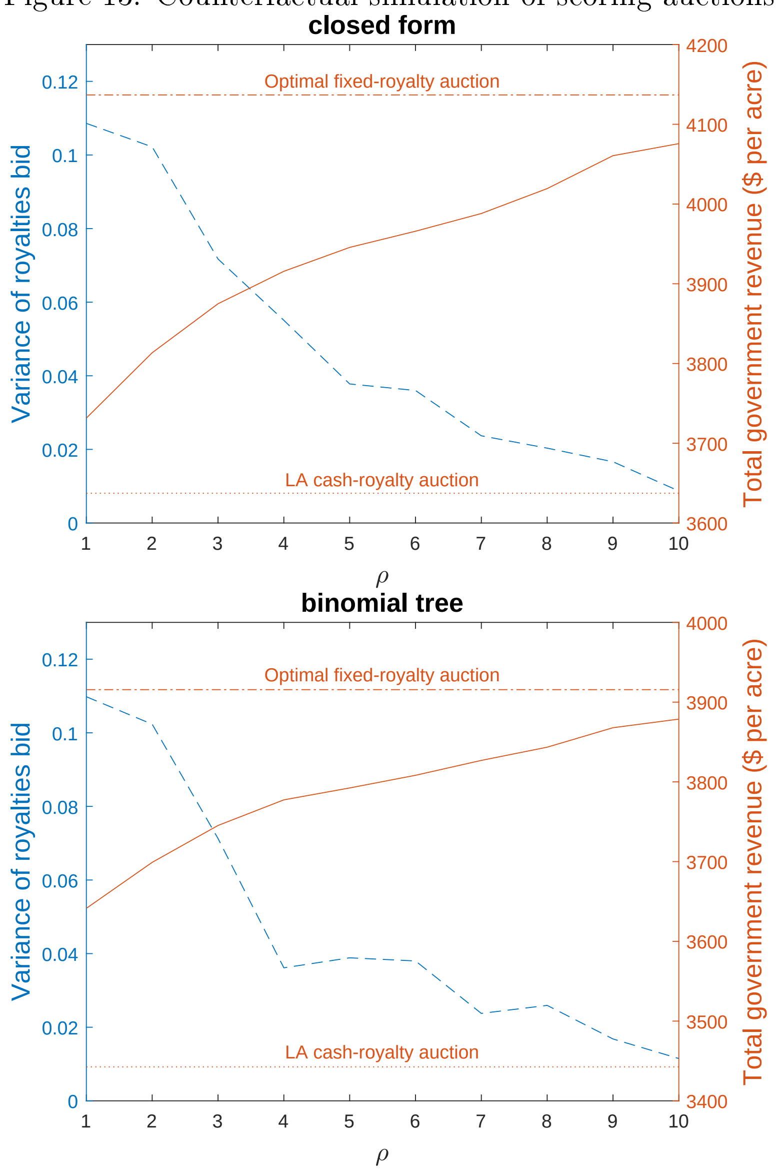

Figure 14: Effect of doubling lease duration

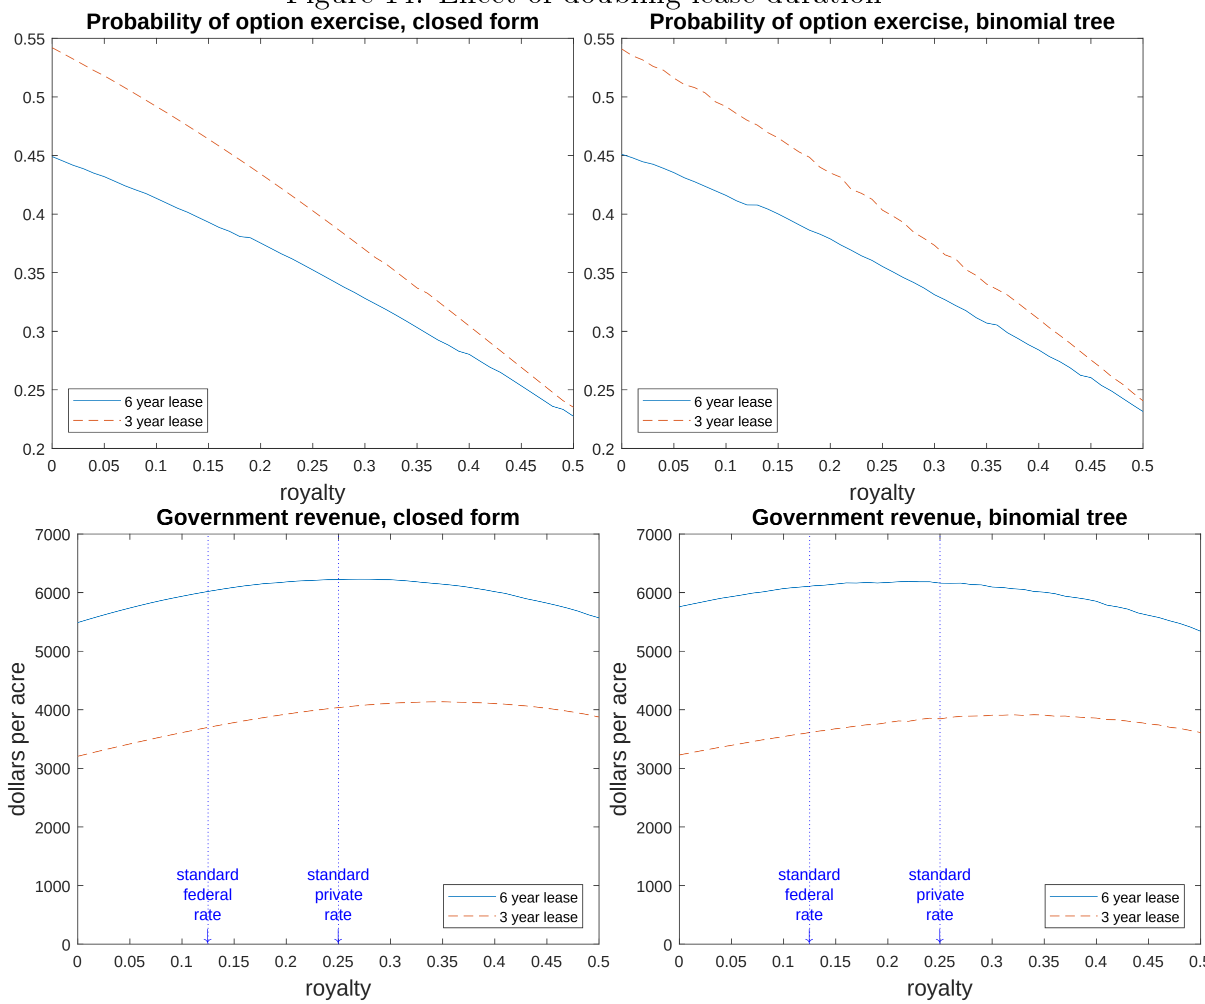

Figure 15: Effect of 20% higher spot price on government revenue

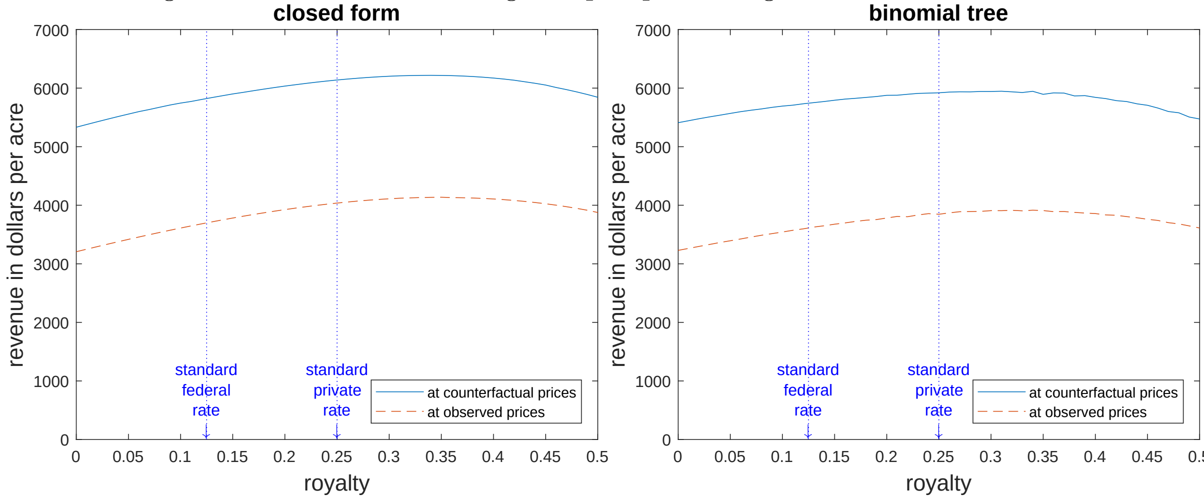

been 20% higher than what they were at the time of each observed auction. The resulting increase in government revenue is much more than proportional; it is 50% or more depending on the royalty. This is because higher oil prices increase not only the royalty dollars conditional on option exercise but also the probability of exercise itself – as shown in Table 2 – as well as option value and hence cash bids.

8 Conclusion

This paper performs a structural analysis of bidding on principal-agent contracts in which both the cash and royalty terms are determined by competitive bidding from potential agents. By using the framework of option values, our model fully accounts for the role of moral hazard: winning agents will engage in production only if they have incentives to do so, and these incentives are affected by the agent's own royalty bid. The analysis is general and flexible, neither depending on a specific scoring function nor assuming optimality for the principal while allowing two-dimensional agent types with a nonparametric distribution.

We show that in Louisiana, the royalty dimension of bidding provides undesirable agents a cheaper currency with which to bid, exacerbating adverse selection in addition to inducing moral hazard. Our counterfactual analysis suggests that reducing the variance of bid royalties increases government revenue, and fixed-royalty auctions obtain more revenue than auctions that allow it as a bid component in addition to cash.

That royalties are better fixed is an important departure from the understanding of multi-attribute auctions that derives from price-quality auctions, in which moral hazard is absent and scoring auctions dominate fixed-quality auctions.

Lastly, auctions of contracts with multidimensional private information is an area where the design space is rich but guiding theory is relatively sparse. It is our hope that this work helps trigger new theoretical and empirical developments in this area.

Appendix

Proof of Proposition 1 First, we show that $V_1(a, \vec{\theta})$ is decreasing in θ_1 and increasing in θ_2. Let θ_k indicate the kth dimension of $\vec{\theta}$. Also define $R \equiv S(1-a)\theta_1$. Then $\frac{\partial}{\partial \theta_k} V_1(a; \vec{\theta}) = \frac{\partial}{\partial \theta_k}\left(\frac{\partial V}{\partial R}\frac{\partial R}{\partial a}\right) = \frac{\partial}{\partial \theta_k}\left(\frac{\partial V}{\partial R}\right)\frac{\partial R}{\partial a} + \frac{\partial}{\partial \theta_k}\left(\frac{\partial R}{\partial a}\right)\frac{\partial V}{\partial R}$. It is well known from the option pricing literature that $\frac{\partial V}{\partial R} = e^{-rt}\Phi(x)$. So $\frac{\partial}{\partial \theta_1}\left(\frac{\partial V}{\partial R}\right) = e^{-rt}\phi(x)\frac{\partial x}{\partial \theta_1} = e^{-rt}\phi(x)\frac{1}{\theta_1\sigma\sqrt{t}}$, and $\frac{\partial}{\partial \theta_2}\left(\frac{\partial V}{\partial R}\right) = e^{-rt}\phi(x)\frac{\partial x}{\partial \theta_2} = e^{-rt}\phi(x)\frac{1}{-\theta_2\sigma\sqrt{t}}$. Meanwhile, $\frac{\partial R}{\partial a} = -S\theta_1$, so $\frac{\partial}{\partial \theta_1}\left(\frac{\partial R}{\partial a}\right) = -S$ and $\frac{\partial}{\partial \theta_2}\left(\frac{\partial R}{\partial a}\right) = 0$. Then, $\frac{\partial}{\partial \theta_1} V_1(a; \vec{\theta}) = -e^{-rt}\phi(x)\frac{1}{\theta_1\sigma\sqrt{t}}S\theta_1 - Se^{-rt}\Phi(x) < 0$, and $\frac{\partial}{\partial \theta_2} V_1(a; \vec{\theta}) = e^{-rt}\phi(x)\frac{1}{\theta_2\sigma\sqrt{t}}S\theta_1 + 0 > 0$.

Next, we show that the choice of royalty $a(p, \vec{\theta})$ is decreasing in θ_1 and increasing in θ_2. We know from (3) that $V_1(a(p, \vec{\theta}), \vec{\theta}) = b_1(a(p, \vec{\theta}), p)\ \forall \vec{\theta}$. Taking a total derivative of both sides with respect to θ_k gives $\frac{\partial}{\partial \theta_k} V_1(a; \vec{\theta}) + \frac{\partial^2 V(a;\vec{\theta})}{\partial a^2}\frac{\partial a(p,\vec{\theta})}{\partial \theta_k}\Big|_{a=a(p,\vec{\theta})} = \frac{\partial^2 b(a,p)}{\partial a^2}\frac{\partial a(p,\vec{\theta})}{\partial \theta_k}\Big|_{a=a(p,\vec{\theta})}$. Rearranging this gives $\left\{\frac{\partial^2 b(a,p)}{\partial a^2} - \frac{\partial^2 V(a;\vec{\theta})}{\partial a^2}\right\}\frac{\partial a(p,\vec{\theta})}{\partial \theta_k}\Big|_{a=a(p,\vec{\theta})} = \frac{\partial}{\partial \theta_k} V_1(a; \vec{\theta})\Big|_{a=a(p,\vec{\theta})}$. By the second-order condition for choosing a, $\frac{\partial^2 b(a,p)}{\partial a^2} - \frac{\partial^2 V(a;\vec{\theta})}{\partial a^2} > 0$. Therefore, $\frac{\partial}{\partial \theta_1} V_1(a; \vec{\theta}) < 0$ from the previous paragraph implies $\frac{\partial a(p,\vec{\theta})}{\partial \theta_1} < 0$, and similarly $\frac{\partial}{\partial \theta_2} V_1(a; \vec{\theta}) > 0$ implies $\frac{\partial a(p,\vec{\theta})}{\partial \theta_2} > 0$.

Proof of Proposition 2 A constructive proof of identification proceeds by proving that there is a unique solution (θ_1, θ_2) to the system (4) and (5). Since x as defined in (2) is a monotonic function of the ratio θ_1/θ_2, we use a change of variables, solving (4) and (5) for θ_1 and x (instead of θ_1 and θ_2) for algebraic convenience. First, as is typically done in solving systems of equations, we use (4) to write θ_1 in terms of x. Second, we plug this expression for θ_1 into (5), yielding an equation with one unknown variable, x. Third, we show that there is a unique solution x to this equation. Finally, we give closed-form expressions for θ_1 and θ_2 as functions of the solution x.

First: Define $R \equiv S(1-a)\theta_1$. Then the derivative $V_1(a,\vec{\theta}) = \frac{\partial V}{\partial R}\frac{\partial R}{\partial a} = -e^{-rt}S\theta_1\Phi(x)$, using a well-known result from option pricing that $\frac{\partial V}{\partial R} = e^{-rt}\Phi(x)$. Substituting this in to the left-hand side of (4), we have $-e^{-rt}S\theta_1\Phi(x) = -P_1(a,b)/P_2(a,b)$. This allows us to write θ_1 in terms of x: $\theta_1 = P_1(a,b)/(P_2(a,b)e^{-rt}S\Phi(x))$. The expression $P_1(a,b)/(P_2(a,b)e^{-rt}S)$ is a positive, known constant; we abbreviate this as C_1, so $\theta_1 = C_1/\Phi(x)$.

Second: Now we plug this expression for θ_1 into the left-hand side of (5) after some rearranging:

$$\begin{aligned} V(a,\vec{\theta}) &\equiv e^{-rt}[S(1-a)\theta_1\Phi(x) - \theta_2\Phi(x-\sigma\sqrt{t})] \\ &= e^{-rt}S(1-a)\theta_1[\Phi(x) - \tfrac{\theta_2}{S(1-a)\theta_1}\Phi(x-\sigma\sqrt{t})] \\ &= \tfrac{e^{-rt}S(1-a)C_1}{\Phi(x)}[\Phi(x) - \tfrac{\theta_2}{S(1-a)\theta_1}\Phi(x-\sigma\sqrt{t})] \\ &= e^{-rt}S(1-a)C_1[1 - \tfrac{\theta_2}{S(1-a)\theta_1}\tfrac{\Phi(x-\sigma\sqrt{t})}{\Phi(x)}] \\ &= e^{-rt}S(1-a)C_1[1 - e^{-\sigma\sqrt{t}x}e^{\sigma^2 t/2}\tfrac{\Phi(x-\sigma\sqrt{t})}{\Phi(x)}], \end{aligned}$$

where the last line comes from plugging in a rearrangement of the definition of x, $\frac{\theta_2}{S(1-a)\theta_1} = \exp(-\sigma\sqrt{t}x + \sigma^2 t/2)$. x is the only unknown in the last line. Now (5) becomes

$$e^{-rt}S(1-a)C_1\left(1 - e^{-\sigma\sqrt{t}x}e^{\sigma^2 t/2}\frac{\Phi(x-\sigma\sqrt{t})}{\Phi(x)}\right) = b + \frac{P(a,b)}{P_2(a,b)}.$$

Collecting x on the left-hand side,

$$e^{-\sigma\sqrt{t}x}\frac{\Phi(x-\sigma\sqrt{t})}{\Phi(x)} = e^{-\sigma^2 t/2}\left(1 - \frac{b + P(a,b)/P_2(a,b)}{e^{-rt}S(1-a)C_1}\right). \tag{6}$$

Third: The right-hand side is known, so we abbreviate it as C_2. If the left-hand side is strictly monotonic in x, there is a unique value of x that satisfies the equation. To check this, we derive the derivative of the left-hand side with respect to x and simplify:

$$\begin{aligned} \frac{\partial}{\partial x}\left(e^{-\sigma\sqrt{t}x}\tfrac{\Phi(x-\sigma\sqrt{t})}{\Phi(x)}\right) &= -\sigma\sqrt{t}e^{-\sigma\sqrt{t}x}\tfrac{\Phi(x-\sigma\sqrt{t})}{\Phi(x)} + e^{-\sigma\sqrt{t}x}\left(\tfrac{\phi(x-\sigma\sqrt{t})\Phi(x)-\Phi(x-\sigma\sqrt{t})\phi(x)}{\Phi(x)^2}\right) \\ &= \left(e^{-\sigma\sqrt{t}x}/\Phi(x)\right)\left(-\sigma\sqrt{t}\Phi(x-\sigma\sqrt{t}) + \phi(x-\sigma\sqrt{t}) - \tfrac{\phi(x)}{\Phi(x)}\Phi(x-\sigma\sqrt{t})\right) \\ &= \left(e^{-\sigma\sqrt{t}x}\phi(x-\sigma\sqrt{t})/\Phi(x)\right)\left(-\sigma\sqrt{t}\tfrac{\Phi(x-\sigma\sqrt{t})}{\phi(x-\sigma\sqrt{t})} + 1 - \tfrac{\phi(x)}{\Phi(x)}\tfrac{\Phi(x-\sigma\sqrt{t})}{\phi(x-\sigma\sqrt{t})}\right) \\ &= \left(e^{-\sigma\sqrt{t}x}\phi(x-\sigma\sqrt{t})/\Phi(x)\right)\left(1 - \tfrac{\Phi(x-\sigma\sqrt{t})}{\phi(x-\sigma\sqrt{t})}\left(\sigma\sqrt{t} + \tfrac{\phi(x)}{\Phi(x)}\right)\right). \end{aligned}$$

By adapting Sampford (1953)'s proof regarding bounds on the derivative of the inverse Mill's ratio, we show in Remark 1 that the derivative of $\frac{\phi(\cdot)}{\Phi(\cdot)} > -1$ for any finite value of the argument. Therefore, $\sigma\sqrt{t} + \frac{\phi(x)}{\Phi(x)} > \frac{\phi(x-\sigma\sqrt{t})}{\Phi(x-\sigma\sqrt{t})} \Rightarrow 1 - \frac{\Phi(x-\sigma\sqrt{t})}{\phi(x-\sigma\sqrt{t})}\left(\sigma\sqrt{t} + \frac{\phi(x)}{\Phi(x)}\right) < 0$ for any value of $\sigma\sqrt{t}$. Then it follows that $\frac{\partial}{\partial x}\left(e^{-\sigma\sqrt{t}x}\frac{\Phi(x-\sigma\sqrt{t})}{\Phi(x)}\right) < 0$, and there is a unique x that satisfies (6). This x, as the only unknown in an equation, is easy to solve for numerically.

Finally, as functions of the solution x, $\theta_1 = P_1(a,b)/(P_2(a,b)e^{-rt}S\Phi(x))$ and $\theta_2 = S(1-a)\theta_1\exp(-\sigma\sqrt{t}x + \sigma^2 t/2)$.

Remark 1. Let $h(x) \equiv \frac{\phi(x)}{\Phi(x)}$. Then $h'(x) > -1$ for $x \in (-\infty, \infty)$.

Proof. We adapt Sampford (1953)'s proof about the derivative of $\frac{\phi(x)}{1-\Phi(x)}$ to prove bounds on the derivative of $\frac{\phi(x)}{\Phi(x)}$. Specifically, consider a standard normal distribution which is top-truncated at x. The variance of this distribution is known to be $1-\frac{x\phi(x)}{\Phi(x)} - \left(\frac{\phi(x)}{\Phi(x)}\right)^2 = 1 - xh(x) - h(x)^2$. Since this is a variance, it must be positive for finite x; $1 - xh(x) - h(x)^2 > 0$. Meanwhile, $h'(x) = -xh(x) - h(x)^2$. Therefore, $1 + h'(x) > 0$, and hence $h'(x) > -1$ for finite x. □

Deriving implied volatility from historical prices of crude oil options We purchase historical prices of crude oil options from the CME Group. For every call option traded in the data, we invert the Black (1976) option pricing equation to back out the expected volatility implied by its price. We take the median of these implied volatilities in each trade month m for option maturity τ (in months) to be the implied volatility in month m of τ-month futures. Implied volatilities derived from 1-month options are noisy, potentially because we observe only the month of option expiration but not the day, so the time left until expiration may not equal an exact month. To address this noise, we adapt Kellogg (2014)'s method of exploiting the term structure of realized volatility. Specifically, we use the term structure to infer the implied volatility of 1-month futures from those of 6-month futures. The term structure of realized volatility that is relevant for each historical month is estimated as follows. For each month, we use daily realized volatilities of oil futures within the surrounding 1-year window to estimate the fixed effects regression $\ln rvol_{t,\tau} = \eta_\tau + \delta_t + \epsilon_{t.\tau}$, where $rvol_{t,\tau}$ is the realized volatility at date t of the τ-month futures contract. The maturity fixed effect η_τ captures the term structure. Finally, the implied volatility $\sigma_{1,m}$ of 1-month futures in month m is inferred from contemporary 6-month implied volatility

$\sigma_{6,m}$ as $\sigma_{1,m} = \sigma_{6,m} \exp(\eta_{1,m} - \eta_{6,m})$.

References

Andreyanov, P. (2018): "Mechanism Choice in Scoring Auctions," http://pandreyanov.com.

Asker, J. and E. Cantillon (2008): "Properties of Scoring Auctions," *The RAND Journal of Economics*, 39, 69–85.

——— (2010): "Procurement When Price and Quality Matter," *The RAND Journal of Economics*, 41, 1–34.

Bhattacharya, V., A. Ordin, and J. W. Roberts (2018): "Bidding and Drilling Under Uncertainty: An Empirical Analysis of Contingent Payment Auctions," https://sites.northwestern.edu/vivek/research/.

Black, F. (1976): "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3, 167–179.

Board, S. (2007): "Selling Options," *Journal of Economic Theory*, 136, 324–340.

Che, Y.-K. (1993): "Design Competition Through Multidimensional Auctions," *The RAND Journal of Economics*, 24, 668–680.

Che, Y.-K. and J. Kim (2010): "Bidding with Securities: Comment," *American Economic Review*, 100, 1929–35.

Cong, L. W. (2019): "Auctions of Real Options," https://www.linwilliamcong.org/research.html.

DeMarzo, P. M., I. Kremer, and A. Skrzypacz (2005): "Bidding with Securities: Auctions and Security Design," *American Economic Review*, 95, 936–959.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68, 525–574.

Hanazono, M., Y. Hirose, J. Nakabayashi, and M. Tsuruoka (2016): "Theory, Identification, and Estimation for Scoring Auctions," http://www.cirje.e.u-tokyo.ac.jp/research/workshops/micro/micropaper16/micro1129.pdf.

Hansen, R. G. (1985): "Auctions with Contingent Payments," *American Economic Review*, 75, 862–865.

Herrnstadt, E., R. Kellogg, and E. Lewis (2019): "Royalties and Deadlines in Oil and Gas Leasing: Theory and Evidence," http://home.uchicago.edu/ kelloggr/KelloggResearch.html.

Hull, J. C. (2017): *Options, Futures, and Other Derivatives*, Pearson, 10th ed.

Kellogg, R. (2014): "The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling," *American Economic Review*, 104, 1698–1734.

Krasnokutskaya, E., K. Song, and X. Tang (2018): "The Role of Quality in Internet Service Markets," https://www.tang-xun.com.

Laffont, J.-J., I. Perrigne, M. Simioni, and Q. Vuong (2018): "Econometrics of Scoring Auctions," .

Laffont, J.-J. and J. Tirole (1987): "Auctioning Incentive Contracts," *Journal of Political Economy*, 95, 921–937.

Larsen, B. and A. L. Zhang (2018): "A Mechanism Design Approach to Identification and Estimation," Working Paper 24837, National Bureau of Economic Research.

Lewis, G. and P. Bajari (2011): "Procurement Contracting with Time Incentives: Theory and Evidence," *The Quarterly Journal of Economics*, 126, 1173–1211.

Li, T., I. Perrigne, and Q. Vuong (2000): "Conditionally Independent Private Information in OCS Wildcat Auctions," *Journal of Econometrics*, 98, 129–161.

McAfee, R. P. and J. McMillan (1987): "Competition for Agency Contracts," *The RAND Journal of Economics*, 296–307.

Nakabayashi, J. (2013): "Small Business Set-Asides in Procurement Auctions: An Empirical Analysis," *Journal of Public Economics*, 100, 28–44.

Sampford, M. R. (1953): "Some Inequalities on Mill's Ratio and Related Functions," *The Annals of Mathematical Statistics*, 24, 130–132.

Sant'Anna, M. (2018): "Empirical Analysis of Scoring Auctions for Oil and Gas Leases," https://marcelosantanna.wordpress.com/.

Skrzypacz, A. (2013): "Auctions with Contingent Payments - An Overview," *International Journal of Industrial Organization*, 31, 666–675.

Takahashi, H. (2018): "Strategic Design under Uncertain Evaluations: Structural Analysis of Design-Build Auctions," *The RAND Journal of Economics*, 49, 594–618.

Yoganarasimhan, H. (2016): "Estimation of Beauty Contest Auctions," *Marketing Science*, 35, 27–54.