

The Performance of Corporate Alliances: Evidence from Oil and Gas Drilling in the Gulf of Mexico

John Beshears^{*}
Harvard University and NBER

July 5, 2013

Abstract

I use data on oil and gas drilling in the Gulf of Mexico to measure how a corporate alliance—a group of firms that jointly develops an offshore tract—performs relative to a solo firm. I employ a regression discontinuity strategy based on bids in first-price sealed-bid auctions for the rights to develop leases. By focusing on leases where one organizational form narrowly outbids the other, I measure drilling outcomes while controlling for the endogenous matching of projects and organizational forms. Solo firm leases are less profitable than alliance leases because alliance members combine their information and expertise.

JEL Classification: G30, G34, D22, D23, D24, L24

Keywords: organizational form, corporate alliances, oil and gas production, lease auctions, regression discontinuity

^{*}jbeshears@hbs.edu, Baker Library 439, Boston, MA 02163, (617)496-6543. For insightful feedback on this research, I thank George Baker, Bo Becker, Daniel Bergstresser, Ryan Bubb, James Choi, Benjamin Esty, Robert Gibbons, Edward Glaeser, Robin Greenwood, Stefan Hunt, Stephanie Hurder, Guido Imbens, Amanda Kowalski, Gregory Lewis, Brigitte Madrian, Sendhil Mullainathan, Joshua Schwartzstein, William Schwert (the editor), Kelly Shue, Ilya Strebulaev, Heidi Williams, an anonymous referee, seminar participants at Boston College Carroll School, the Federal Reserve Board, Harvard Business School, London School of Economics, RAND Corporation, Stanford Graduate School of Business, UC Berkeley Haas School, University of Chicago Booth School, University of Pennsylvania Wharton School, and Yale School of Management, and particularly my dissertation committee: Malcolm Baker, Max Bazerman, David Laibson, and Jeremy Stein. James Jenkins and Russell Stauffer were invaluable guides to the oil and gas industry. I am grateful to Marianne Bertrand and Sendhil Mullainathan for sharing their data. Kyle Chauvin, Jessica Matthews, Vanessa Schneider, and Brian Son provided excellent research assistance. I acknowledge financial support from a National Science Foundation Graduate Research Fellowship, Harvard Business School, and Stanford Graduate School of Business.

1. Introduction

In this paper, I address a question in the tradition of Coase (1937): how does organizational form affect investment performance? Mergers, acquisitions, joint ventures, and corporate alliances may combine firms' capabilities to create synergies, alter the nature of agency relationships, or reallocate decision rights, and a central issue in corporate finance is to understand the economic implications of these changes in firm boundaries. To shed light on this issue, I study the case of corporate alliances developing oil and gas leases in the Gulf of Mexico and compare their performance to the performance of solo firms.

Corporate alliances—groups of firms that enter long-term contractual arrangements for the purposes of jointly pursuing investment projects—are a ubiquitous organizational structure in many sectors of the economy. According to Baker, Gibbons, and Murphy (2008), the top firms in the pharmaceutical and biotechnology industries are entangled in a complex network of alliances, with more than a thousand firms directly connected via alliances to a core group of a few dozen industry leaders. Other examples of alliances abound in industries such as telecommunications, medical devices, and software.

In the case of Gulf of Mexico oil and gas leases auctioned by the U.S. federal government, nearly two-thirds of the more than \$300 billion in bids submitted between 1954 and 2007 were from alliances.¹ Oil and gas drilling in the Gulf of Mexico has been overseen by the U.S. Department of the Interior since the 1950s. Firms can participate in first-price sealed-bid auctions for the right to extract the mineral resources in offshore tracts, and the winning bidders pay royalties to the Department of the Interior on any

¹ For this calculation, bid values are adjusted to 2007 dollars.

production from the leases. Firms are permitted to form alliances for the purposes of jointly acquiring and developing leases, or they can simply submit bids as individual entities.

This setting provides several advantages for studying the investment outcomes of corporate alliances. First, offshore drilling is an activity where sophisticated planning and analysis are necessary for success, so some of the features of alliances that have been posited as advantageous properties, such as their ability to combine the information and expertise of their members (see Doz and Hamel, 1998), might be expected to affect investment performance.² Second, the Department of the Interior has assembled detailed and comprehensive data on the characteristics and productivity of individual drilling projects in the Gulf of Mexico, allowing me to construct estimates of the costs and revenues from offshore activity at the borehole level and at the lease level.³ Third, because leases are allocated in first-price sealed-bid auctions, I can employ a regression discontinuity strategy to compare the outcomes of alliances and solo firms while controlling for lease heterogeneity. It is likely that alliances and solo firms are endogenously matched to leases with different characteristics, so it could be misleading to examine differences in outcomes between the set of all leases acquired by alliances and the set of all leases acquired by solo firms. The regression discontinuity strategy addresses this difficulty by focusing on leases where an alliance narrowly outbids a solo firm or a solo firm narrowly outbids an alliance. If each of the bidders is choosing its bid

² Other rationales for alliance formation, such as improved access to consumer markets, are less relevant in this setting.

³ In this regard, my analysis builds upon previous work that uses offshore oil and gas drilling as a “laboratory” for empirically examining the decision-making of firms. Hendricks, Porter, and Boudreau (1987) launched a series of papers that use the rich data from the Gulf of Mexico to study firms’ strategic behavior. My calculations of drilling profitability are largely based on their methodology. Bertrand and Mullainathan (2005) use the same data to evaluate theories of the relationship between firm investment and cash flow.

based on its own signal regarding the profitability of the lease, the situation in which the alliance narrowly outbids the solo firm and the situation in which the bids are slightly perturbed such that the solo firm narrowly outbids the alliance reflect nearly identical signal configurations, so the lease characteristics and the severity of the winner's curse are also nearly identical in the two situations. Intuitively, as the highest alliance bid and the highest solo firm bid become closer in value, the allocation of the lease to an alliance or a solo firm approaches "random assignment," and it is therefore possible to compare alliance and solo firm investments while controlling for lease characteristics.⁴

Using this regression discontinuity approach to examine leases auctioned between 1954 and 1975, I find that boreholes drilled by solo firms are less profitable than those drilled by alliances by approximately \$1.6 million (in 1980 dollars) per borehole. Furthermore, leases acquired by solo firms have a lower net present value than leases acquired by alliances. The magnitude of this effect is approximately \$31 million (in 1980 dollars) per lease.

The evidence suggests that alliances achieve better drilling outcomes than solo firms. It is important to note, however, that the superior performance of alliances does not imply that alliances and solo firms cannot coexist in a long-run industry equilibrium. Even if alliances have better performance than solo firms according to my measures of drilling profitability, the formation and maintenance of alliances incur offsetting costs that are not directly associated with drilling. These costs may be overhead expenses such as legal fees, or they may reflect managerial objectives that diverge from the interests of

⁴ Section 4.1 provides a detailed description of the estimation procedure along with a simple model that develops the intuition behind the identification strategy. This empirical approach differs from another common approach to the study of organizational form, which examines how organizational form varies with project characteristics (see, for example, Monteverde and Teece, 1982; Masten, 1984; and Baker and Hubbard, 2003, 2004).

firm owners, such as a manager's preference to pursue lease development outside of an alliance in order to preserve decision-making autonomy. The estimated differences in drilling performance between alliances and solo firms are large—approximately half the standard deviations of the outcome variables—so if equilibrium holds, the results imply that the non-drilling offsetting costs must also be substantial. From this perspective, the results illuminate the relative advantages and disadvantages of alliances as an organizational form.

To explore the potential sources of the difference between alliance and solo firm drilling outcomes, I begin by breaking down drilling profitability into the component due to drilling expenditures and the component due to operating profits from oil and gas production subsequent to drilling. Differences in profitability between alliances and solo firms are driven by the latter component. If anything, alliances expend more resources on drilling than solo firms, a factor pushing in the opposite direction of the main results.

Furthermore, I sort alliances into categories based on how many of their member firms have a high degree of previous experience owning leases in the geographic area around the lease to be developed. Relative to alliances with two or more high-experience member firms, alliances with zero or one high-experience member firm have drilling outcomes that are more similar to solo firm drilling outcomes.

Collectively, these findings are consistent with the hypothesis that alliances achieve superior drilling outcomes by combining the information and expertise of their member firms (Doz and Hamel, 1998). The members of an alliance, by virtue of their previous experience drilling in the Gulf of Mexico or elsewhere, may each possess distinct pieces of information that can be brought to bear on the question of where to drill

within a lease. For example, if an alliance is deciding where to drill in a particular type of geological structure, the alliance members may each have experience dealing with similar structures, and they may each possess unique insights into the best way to proceed with drilling. By aggregating that information, alliances may be able to choose drilling locations that are more productive than the locations chosen by solo firms.

The hypothesis that alliances are able to achieve better outcomes by aggregating the information and expertise of their members is natural and intuitive, but there are other possible explanations for the superior performance of alliances. First, it may be that alliances help resolve agency problems that exist at the level of individual firms. However, the finding that alliances have higher drilling expenditures than solo firms suggests that alliances do not overcome a tendency for solo firm managers to be too aggressive. At the same time, it turns out that alliances are slightly less likely to commence drilling on a lease than solo firms, suggesting that alliances do not counteract a general tendency for solo firm managers to be too conservative in their drilling strategies. Furthermore, in quantile regressions, the drilling strategies of alliances do not appear to have greater downside risk than those of solo firms, so it does not seem that alliances encourage managers to shift towards riskier strategies with higher expected returns. Subtler forms of agency conflicts may be at play, but the simplest versions of agency problems are not consistent with the data.

Second, alliances may have better performance because they are less aggressive in their bidding strategies. If alliances and solo firms have similar valuations for leases but alliances submit bids that are well below those valuations while solo firms submit bids that are only somewhat below those valuations, alliances will have better drilling

outcomes than solo firms conditional on one organizational form narrowly outbidding the other. Looking at the bid distributions for leases close to the discontinuity and their neighbors, however, I find that alliances are more aggressive than solo firms in their bidding strategies.

Third, it is possible that high-quality firms are more likely to enter alliances, perhaps because it is easier for them to attract alliance partners. However, when I incorporate different forms of firm fixed effects into the regression analysis, the main results are not altered, suggesting that the selection of firms into organizational forms is not driving the superior outcomes of alliances.

Fourth, alliances and solo firms may have differences in discount rates or financial resources. Introducing control variables capturing leverage and available resources into the regression framework does not alter the main findings, so these factors do not appear to explain the results.

Fifth, systematic measurement error in drilling costs could help account for the pattern of alliances exhibiting superior drilling performance. However, the main results are qualitatively similar under reasonable alternative assumptions regarding drilling costs.

Sixth, alliances may create stronger incentives for managerial investment in information relevant to the development of a lease. It is difficult to provide evidence that supports or refutes this hypothesis, but this hypothesis is similar to the explanation that alliances bring combined information and expertise to bear on a given lease.

These results are not meant to strictly rule out the role of agency issues, bidding strategies, selection, discount rates, financial resources, or differences in drilling costs as factors contributing to the finding that alliances have superior performance relative to

solo firms. However, the pattern of evidence points towards alliances' ability to combine the information and expertise of their members as an important source of their advantage.

This research contributes to the literature on organizational design and the boundaries of the firm. Previous theoretical work that adopts the property rights approach to the theory of the firm (Grossman and Hart, 1986; Hart and Moore, 1990) argues that hybrid organizational forms such as alliances can optimally balance the incentives of firms to contribute to a joint production process (Robinson, 2008; Fulghieri and Sevilir, 2009; Hackbarth, Mathews, and Robinson, 2012). The results in this paper provide support for the premise that alliances can combine inputs from multiple firms to generate productivity advantages relative to other organizational forms. Moreover, the results shed light on the *types* of inputs that alliances are designed to encourage. As emphasized in the theoretical work of Doz and Hamel (1998), Mathews (2006), and Habib and Mella-Barral (2007, 2013), the sharing of information and expertise is a central motive for alliance formation.

Previous empirical research on alliances has concluded that firms exhibit improved performance when they are involved in alliances (McConnell and Nantell, 1985; Mitchell and Singh, 1996; Chan et al., 1997; Anand and Khanna, 2000; Stuart, 2000).⁵ However, there is less research that directly studies the performance of alliances themselves relative to the performance of other organizational forms. Mowery, Oxley, and Silverman (1996) and Gomes-Casseres, Hagedoorn, and Jaffe (2006) find that

⁵ There is also a body of work that examines the determinants of alliance formation and structure. See, for example, Gulati (1995a, b), Lerner and Merges (1998), Elfenbein and Lerner (2003), Lerner, Shane, and Tsai (2003), Desai, Foley, and Hines (2004), Robinson and Stuart (2007a, b), and Lerner and Malmendier (2010). Of course, there is a vast literature on the determinants and effects of vertical integration (for a review, see Lafontaine and Slade, 2007), which I treat as distinct because the alliances I examine are primarily horizontal in nature.

alliances affect the pattern of patent cross-citations between allying firms, but they do not study the impact on financial performance. Kent (1991) examines alliance performance in the Gulf of Mexico but concludes that alliances do not have better outcomes than solo firms, which is what I also find if I do not use the regression discontinuity methodology to control for the endogenous assignment of leases. In a study of the pharmaceutical industry, Guedj (2005) finds that drug candidates developed jointly by a large firm and a small firm are more likely to be continued but are ultimately less successful than projects conducted solely by a large firm. Interestingly, my results are the inverse of his—at the regression discontinuity, alliances are slightly less likely to pursue drilling, and their drilling is more profitable than that of solo firms. The difference in results may be due to the conflicting interests of the large firm and small firm alliance partners in the Guedj study. Guedj argues that a given small firm in his sample has a greater interest in continuing a drug development project than the allied large firm, since the project often represents one of the small firm’s only investment opportunities. In my sample of oil and gas firms, on the other hand, the interests of alliance members are better aligned, and the advantageous features of alliances, such as their capacity for combining information and expertise, can drive the performance outcomes.⁶

Finally, this research is related to previous work on oil and gas firm activity in the Gulf of Mexico. Hendricks and Porter (1992) and Hendricks, Porter, and Tan (2008) study alliances in the Gulf of Mexico, but they focus primarily on joint bidding in lease auctions, while I emphasize the implications of alliances for drilling outcomes after the

⁶ I focus here on the literature studying alliances, but there is a broader empirical literature examining the connection between investment decisions and other aspects of organizational structure, such as the separability of residual and fixed claims (Esty, 1997), bureaucratic complexity and room for opportunism (Sampson, 2004), firm size (Berger et al., 2005), firm scope (Guedj and Scharfstein, 2005; Seru, forthcoming), and firm status as public or private (Sheen, 2009).

auctions conclude. Hendricks and Porter (1996) do analyze drilling decisions, although they do not devote their attention to comparing the drilling decisions of alliances and solo firms.

The remainder of the paper proceeds as follows. Section 2 provides information on the institutional environment in the Gulf of Mexico, and Section 3 describes the data set. Section 4 explains the empirical methodology, develops the intuition behind the identification strategy using an illustrative model, and presents the main results on the superior drilling outcomes of alliances relative to solo firms. In Section 5, I discuss possible explanations for the main results, and Section 6 concludes.

2. Institutional details of oil and gas drilling in the Gulf of Mexico

Offshore drilling for oil and gas has occurred in the Gulf of Mexico since the 1940s. The Outer Continental Shelf Lands Act, passed by the United States Congress in 1953, gave the Department of the Interior responsibility for conducting lease auctions, collecting royalties, and generally administrating all activity undertaken in federal waters. My discussion of the institutional features governing drilling in this area closely follows the explanations of Hendricks, Porter, and Boudreau (1987), Hendricks and Porter (1992), and Hendricks, Porter, and Tan (2008).

Before conducting a sale of Gulf of Mexico leases, the Minerals Management Service (MMS)⁷ of the Department of the Interior consults with potential bidders regarding which tracts would garner interest. MMS then designates certain tracts to be

⁷ Due to a reorganization in response to the Deepwater Horizon oil spill in 2010, the Bureau of Ocean Energy Management and the Bureau of Safety and Environmental Enforcement have replaced the Minerals Management Service. I discuss the now defunct institution for the sake of historical accuracy and for the sake of continuity with previous literature on Gulf of Mexico leases.

included in the lease sale. The lease terms are fixed, with a typical lease having an area of 5,000 acres and a royalty rate of 16 $\frac{2}{3}$ %. Each lease is sold in a separate auction. Bidders compete in an auction by offering a “bonus,” which is an up-front payment to MMS due at the time of the auction if the bid is accepted. The auctions use a first-price sealed-bid format. The highest bid wins the lease, although MMS reserves the right to refrain from issuing a lease when it deems the highest bid too low. MMS simultaneously conducts the auctions for tens or hundreds of leases in a single sale, and sales take place at a rate of approximately one or two per year. Firms cannot conduct drilling on a lease before the sale occurs, but they are permitted to perform seismic surveys on a lease in preparation for submitting a bid.

Firms are allowed to form alliances for the purposes of bidding jointly on any given lease. Members of an alliance for one lease are permitted to bid in auctions for other leases as a solo firm or as part of a different alliance. In practice, the contracts signed by members of an alliance bind them to bid jointly on a handful of leases that are part of a single sale and are within a single geographic area. At the end of 1975, because of concerns that joint bidding reduced the competitiveness of lease auctions, a new federal law and new Department of the Interior regulations forbade the eight largest oil firms from forming alliances with each other. Since these restrictions changed the nature of alliances operating in the Gulf of Mexico (see Hendricks and Porter, 1992), my empirical analysis is limited to leases that were auctioned during the period 1954–1975.

If a firm or group of firms is granted a lease, there is a fixed term, generally five years, during which the firm or group of firms has the exclusive option to drill on the lease. If production of minerals begins on the lease, the lease term is automatically

extended until production ceases. Otherwise, the lease reverts to MMS at the end of the initial term. MMS collects a fixed fraction of all production revenue from a lease as a royalty payment.

3. Construction and summary of the data set

3.1. Data set construction

Data on Gulf of Mexico lease auctions, drilling, and mineral production are from the Minerals Management Service of the Department of the Interior. There were 2,636 offshore Louisiana and Texas leases that were put up for sale between 1954, the first year for which information is available, and 1975, the last year before restrictions on the formation of certain alliances took effect. I limit the sample to leases sold using the auction format described in Section 2,⁸ and I drop the 296 leases that were put up for sale but not issued, primarily as a result of the highest bid being rejected by the Minerals Management Service as inadequate.

For each lease, the auction data include the identities of all bidding companies and the amounts of all bids. In the case of an alliance bidder, the data also include the percent of the bid allocated to each member of the alliance.⁹ Bidding records for eight leases were incomplete or unclear, so these leases are excluded from the sample.¹⁰

⁸ This restriction eliminates eight leases that were sold in “royalty bid” auctions, where participants submitted bids indicating the percentage royalty they were willing to pay on production revenue.

⁹ In some cases, a company appears in the data under multiple names, for instance when a parent company and its subsidiary or affiliate are both involved in bidding. I use name matching, Internet searches of company histories, and the *Directory of Corporate Affiliations* to combine the different occurrences of a company and treat them as a single entity, pooling their ownership shares in situations where they are involved in the same bid.

¹⁰ Five of these leases were auctioned in 1954; one was auctioned in 1968; and the remaining two were auctioned in 1970.

A bidder is considered a solo firm if it consists of a single company that is wholly responsible for the bid. A bidder is considered an alliance if it consists of two or more companies that are jointly responsible for the bid, and if none of those companies is allocated more than half of the bid. The remaining bidders are consortia where one of the member companies is responsible for strictly greater than 50% of the bid. These bidders may have some of the properties of solo firms, since the company with the majority stake may be able to exert particular influence over the consortium's decisions. At the same time, these bidders may have some of the properties of alliances, since financial risks are not borne solely by the company with the majority stake. Because the classification of these consortia is ambiguous, I remove from the data set all of their bids as well as the 155 leases where they were the winning bidder.¹¹

To study the drilling outcomes of solo firms and alliances while controlling for tract characteristics, the regression discontinuity strategy relies on comparing tracts where a solo firm narrowly outbid an alliance to tracts where an alliance narrowly outbid a solo firm. Because this strategy requires both a solo firm and an alliance to participate in the bidding, my primary analysis sample is limited to leases where this is the case. Restricting the sample in this way eliminates 1,095 leases.

To quantify the closeness of the bidding, I construct the variable *BIDRATIO* as follows. For each tract where a solo firm won the bidding but an alliance also participated, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm's bid. For a tract where an alliance won but a solo firm participated, *BIDRATIO* is defined as the highest solo firm bid divided by the winning alliance's bid. When

¹¹ I have experimented with classifying these ambiguous consortia as solo firms or as alliances. As might be expected given their limited presence in the data, the main results are largely unaffected.

calculating this variable, I ignore any bid submitted by a consortium where one member was responsible for more than half of the bid.

In addition to *BIDRATIO*, other lease-level variables I observe include the start date of the lease and changes in lease ownership after the auction.¹² I also observe the geographic locations of leases. More specifically, I observe the Minerals Management Service (MMS) “area” and “block” in which each lease is located. MMS has partitioned the Gulf of Mexico into “blocks,” which are typically rectangular geographic regions that are approximately 5,000 acres in size. Many but not all leases cover a geographic region that exactly corresponds to one block. “Blocks” are grouped into “areas,” which are clusters of approximately 50 to 250 contiguous “blocks.” I am unable to identify the location of one lease, so I drop it from the sample. In addition, I have insufficient data to use the methodology described below to calculate an estimate of drilling costs for boreholes on three remaining leases, which I also drop from the sample. I am left with a primary analysis data set consisting of 1,070 leases.

I use data on drilling activity and mineral production to calculate ex post economic returns from investment at the borehole level and at the lease level. When analyzing economic returns at the borehole level, I restrict attention to boreholes drilled on a lease within five years of the lease start date. All leases in my sample had initial terms of at least five years,¹³ but a lease term was automatically extended if production began on the lease. Limiting the analysis to boreholes drilled within five years of the lease start date ensures that all observations are drawn from comparable timeframes. When analyzing economic returns at the lease level, I consider all boreholes drilled

¹² In cases where I know a company was acquired, I treat the event as an ownership transfer of the company’s leases to the acquirer.

¹³ Seven leases had initial terms of ten years, and all others had initial terms of five years.

within 20 years of the lease start date, since including these boreholes better reflects the total economic value generated from a lease.

My methodology for calculating economic returns largely follows that of Hendricks, Porter, and Boudreau (1987). Differences are noted in the Appendix, which also provides details on all calculations.

There were 5,833 (10,415) boreholes drilled on the 1,070 leases in my sample within five (20) years of the lease start date. I observe each borehole's total vertical depth, the distance along its axis from the drilling rig to the bottom, the water depth at its location, the latitude and longitude of its bottom, the latitude and longitude of the drilling rig, the elevation of the drilling rig above water level, and the spud date (the date on which drilling commenced). I use this information to construct the variable *BHCOST*, which is an estimate of the cost of drilling and equipping a borehole. The American Petroleum Institute (API) has conducted an annual survey on drilling costs since the 1950s, with costs tabulated by region, well type, and well depth. I calculate the number of feet drilled for each borehole and multiply it by the cost per foot from the API survey for the appropriate year, region (offshore Louisiana or Texas), well type (oil, gas, or unproductive), and well depth. In the case of boreholes that ultimately produced oil and gas, I further adjust the drilling cost by multiplying it by a factor that accounts for the cost of additional lease equipment. This factor, which varies by year, is based on data from the API survey, from the Census Bureau's Annual Survey of Oil and Gas, and from the Energy Information Administration (EIA). Finally, all costs are converted to 1980

dollars using the gross domestic product (GDP) implicit price deflator index from the Bureau of Economic Analysis.¹⁴

It is important to note that the estimates of drilling costs are the component of the estimated economic returns to drilling that are most subject to measurement error. The calculation of the number of feet drilled for a given borehole is quite precise, but the number of feet drilled is then multiplied by a cost per foot of drilling that is based on the average cost per foot among similar boreholes, where similarity is judged by time period, geographic location, type of well, and well depth. The actual cost of drilling a given borehole could be higher or lower than this estimate. I examine the sensitivity of the empirical results to the possibility of systematic measurement error in the cost of drilling in Section 5.6.

I calculate the operating profit from each borehole as follows. I multiply monthly oil production by the average offshore Louisiana wellhead crude oil price for the year, converted to 1980 dollars, and I multiply monthly gas production by the average offshore Louisiana wellhead natural gas price for the year, also converted to 1980 dollars.¹⁵ The sum of these two quantities is monthly revenue. I subtract off a fraction of monthly revenue to account for royalty payments to the Minerals Management Service, and I subtract off a further fraction of monthly revenue to account for well operating costs. This second fraction is based on data from the API survey, from the Census Bureau, and from the EIA. To calculate a net present value as of the spud date for this stream of operating profits, I consider the sequence of monthly operating profits starting at the spud

¹⁴ The data span the years 1954 to 2006, so 1980 represents the middle of the relevant time period.

¹⁵ Louisiana wellhead prices are not available for years prior to 1960. However, oil and gas prices were quite stable in the 1950s. My calculations assume that the real prices of oil and gas were constant over the period 1954–1960.

date and extending through 25 years or through the year 2006, whichever is earlier. For subsequent years, I take the level of production from the earlier of the 25th year and the year 2006, and I assume that it declines at a 25% annual rate in perpetuity. Operating profits from this final production profile are calculated using oil prices, gas prices, and operating cost assumptions from the end of the 25th year or from the year 2006. The variable *BHOP* is equal to the sum of all operating profit attributed to the borehole, discounted to the spud date using a 5% annual rate.

Finally, the net present value of a borehole, *BHNPV*, is simply the present discounted value of operating profits *BHOP* minus drilling and equipping costs *BHCOST*. For a given lease, I define *MEANBHNPV* as the mean of *BHNPV* taken over all boreholes drilled within five years of the lease start date. To calculate the net present value of a lease, *LSNPV*, I sum *BHNPV* over all boreholes drilled on the lease within 20 years of the lease start date, discounting back to the lease start date at a 5% annual rate. I then subtract off the bonus paid to acquire the lease. *MEANBHNPV* and *LSNPV* are each winsorized at the 1st and 99th percentiles.

3.2. Data set summary

To provide context for the analysis of alliance and solo firm outcomes, I first summarize the bidding behavior in the primary analysis sample. Table 1 lists the ten most active firms, as determined by the number of dollars they contributed to winning bids in the sample. The table quantifies each firm's solo activity and alliance activity by reporting the number and dollar amount of winning bids of both types.

There were some major firms that participated in lease auctions almost exclusively as solo bidders. For example, 82 out of the 92 winning bids submitted by

Shell were solo bids, and these bids constituted 82.4% of the \$878 million that Shell contributed to winning bids in the sample. On the other hand, some firms predominately participated as part of an alliance. Getty contributed \$724 million to 112 winning joint bids but paid only \$34 million for two winning solo bids. However, most firms engaged in a balanced mix of solo and alliance bidding. Exxon, the most active firm in the sample, spent \$1.04 billion on 61 winning solo bids and \$482 million on 27 winning alliance bids. Overall, 47.4% of the winning bids in the sample were solo bids. These bids represented 32.7% of the \$17.51 billion devoted to winning bids, reflecting the fact that solo bids tended to be smaller than alliance bids.

Table 2 reports the distributions of lease-level variables. Most leases had an area of approximately 5,000 acres, although some leases were smaller. The distribution of *BIDRATIO*, which is the key variable for the regression discontinuity strategy, has somewhat higher density at values close to zero than at values close to one, but there is still a substantial number of leases where the highest alliance bid and highest solo firm bid were quite similar. The mean borehole profitability variable *MEANBHNPV* and the lease net present value variable *LSNPV* both have distributions with long right-hand tails due to the nature of offshore drilling. The costs of drilling are large and vary with well depth and other technological factors, but most of the variability in the profitability of drilling is driven by the presence and extent of mineral resources. Drilling occasionally results in the discovery of a major reservoir. This outcome is rare but highly profitable, and it compensates for the many drilling projects that do not ultimately produce minerals in meaningful quantities. Because these outcomes in the right-hand tail are an important aspect of the economics of offshore drilling, I largely leave them intact in the data. At

the same time, I do not want extreme outliers to exert too much influence on my regression analyses, which is why *MEANBHNPV* and *LSNPV* are winsorized at the 1st and 99th percentiles. Note that long right-hand tails are also apparent in the distributions of winning bids, the number of boreholes drilled within five years of the lease start date, and the number of boreholes drilled within 20 years of the lease start date.

Table 3 provides summary statistics at the borehole level. The number of feet drilled for a borehole ranges from several hundred to several thousand feet, and *BHCOST* reflects this distribution because I calculate the drilling and equipping cost by multiplying the number of feet drilled by an estimate of cost per foot. Borehole operating profit *BHOP* and borehole net present value *BHNPV* exhibit long right-hand tails. Indeed, the median borehole in my sample has negative net present value. *BHCOST*, *BHOP*, and *BHNPV* are winsorized at the 1st and 99th percentiles for the purposes of this table.

4. Main results comparing the drilling outcomes of alliances and solo firms

In this section, I first explain the regression discontinuity design and the assumptions that underlie it, using an illustrative model to develop the intuition behind the identification strategy. I then describe the implementation of the strategy. Finally, I report the main regression discontinuity results that compare borehole profitability and lease profitability for alliances and solo firms.

4.1. Interpretation of the regression discontinuity strategy

To compare the drilling outcomes of alliances and solo firms while controlling for lease heterogeneity, I use a regression discontinuity design that examines leases for which a solo firm narrowly outbid an alliance or an alliance narrowly outbid a solo firm.

For these leases, the variable *BIDRATIO* has a value close to one. Recall that *BIDRATIO* is defined as the highest alliance bid divided by the winning solo firm bid when a solo firm wins the bidding and is defined as the highest solo firm bid divided by the winning alliance bid when an alliance wins the bidding.

For the purposes of discussing the empirical design, it is convenient to define the variable R^* as follows:

$$R^* = \begin{cases} \textit{BIDRATIO} & \text{if a solo firm wins the auction} \\ 2 - \textit{BIDRATIO} & \text{if an alliance wins the auction} \end{cases} . \quad (1)$$

As R^* approaches one from below, the highest alliance bid approaches the winning solo firm bid, and as R^* approaches one from above, the highest solo firm bid approaches the winning alliance bid. I am interested in comparing drilling outcomes as R^* approaches one from below to drilling outcomes as R^* approaches one from above, since the allocation of a lease to a solo firm or to an alliance changes discontinuously at $R^* = 1$ while the mix of lease types is expected to change continuously at $R^* = 1$. Thus, differences in outcomes at the discontinuity can be attributed to the assignment of the lease to a solo firm or to an alliance, instead of attributed to heterogeneity in lease types.

The following simple model, which is an example of the “wallet game” of Bulow and Klemperer (2002), illustrates the intuition behind the regression discontinuity strategy. The specific parameterization of the model described here is a trivial extension of the parameterization studied by Avery and Kagel (1997), so I refer the reader to their appendix for a proof of the claims regarding equilibrium bidding strategies.

Two bidders compete in a first-price sealed-bid auction for a single lease. Bidder one privately observes a signal S_1 regarding the value of the lease, and bidder two

privately observes a signal S_2 regarding the value of the lease, with the two signals independently distributed uniformly on the unit interval. Bidder one receives a payoff of $fS_1 + (1-f)S_2 + v$ minus the value of the winning bid for winning the lease auction and developing the lease, while bidder two receives a payoff of $fS_1 + (1-f)S_2$ minus the value of the winning bid for winning the lease auction and developing the lease, where $\frac{1}{2} < f < 1$ and $v > 0$. Note that bidder one has two advantages relative to bidder two in this setting: bidder one's signal is more informative regarding the common component of the value of the lease ($f > \frac{1}{2}$), and bidder one receives a larger payoff from winning the auction, conditional on the signals and the winning bid ($v > 0$). Bidder one can be thought of as an alliance, and bidder two can be thought of as a solo firm.

Given bidder two's bidding strategy $b_2(s_2)$ and a signal realization of s_1 , bidder one chooses a bid b_1 that maximizes $\int_0^{b_2^{-1}(b_1)} (fS_1 + (1-f)S_2 + v - b_1) dS_2$. Similarly, given bidder one's bidding strategy $b_1(s_1)$ and a signal realization of s_2 , bidder two chooses a bid b_2 that maximizes $\int_0^{b_1^{-1}(b_2)} (fS_1 + (1-f)S_2 - b_2) dS_1$. Avery and Kagel (1997) show that for a sufficiently small value of v there is an equilibrium of the game featuring bidding strategies of the form

$$b_1(s_1) = a_1 s_1 + c_1 \quad (2)$$

$$b_2(s_2) = a_2 s_2 + c_2 \quad (3)$$

where

$$a_1 = \left(\frac{1}{4} - \frac{1}{2}v\right) + \sqrt{\left(\frac{1}{4} + \frac{1}{2}v\right)^2 - \frac{1}{2}(1-f)v} \quad (4)$$

$$a_2 = \left(\frac{1}{4} + \frac{1}{2}v\right) + \sqrt{\left(\frac{1}{4} + \frac{1}{2}v\right)^2 - \frac{1}{2}(1-f)v} \quad (5)$$

$$c_1 = 2\left[\left(\frac{1}{4} + \frac{1}{2}v\right) - \sqrt{\left(\frac{1}{4} + \frac{1}{2}v\right)^2 - \frac{1}{2}(1-f)v}\right] \quad (6)$$

$$c_2 = 2\left[\frac{1}{4} - \sqrt{\left(\frac{1}{4} + \frac{1}{2}v\right)^2 - \frac{1}{2}(1-f)v}\right]. \quad (7)$$

The intercept of bidder one's bid function is greater than the intercept of bidder two's bid function, but the slope of bidder two's bid function is greater than the slope of bidder one's bid function. The two bid functions generate the same bid for the highest possible signal realization.¹⁶

In this model, the bidding strategies are asymmetric. The different bidders win the lease auction for different lease types, where a lease's type is defined by the pair of signal realizations, so a simple comparison of outcomes for leases won by bidder one (an alliance) and leases won by bidder two (a solo firm) would capture both differences driven by the productivity of alliances versus solo firms and differences driven by heterogeneity in lease types.

The idea behind the regression discontinuity strategy is to control for heterogeneity in lease types by comparing outcomes for leases won by the alliance and leases won by the solo firm *conditional on* the two bidders submitting (nearly) the same bid, that is, conditional on $R^* = 1$. Within the model, the difference between the expected payoff for the alliance and the expected payoff for the solo firm, conditional on $R^* = 1$, is exactly equal to v .¹⁷ To see the reason for this claim, consider the case in which both the

¹⁶ The requirement that bids be nonnegative implies that the bidding strategy specified above for bidder two does not literally apply. When the equations above call for bidder two to submit a negative bid, bidder two actually submits a bid of zero. This adjustment does not affect the conclusions from the model because it does not change bidder one's optimal strategy and because bidder two does not win the auction for low signal realizations anyway.

¹⁷ Note that the measure of borehole profitability *MEANBHNPV* and the measure of lease profitability *LSNPV* described in Section 3.1 do not map directly to the parameter v in the model, as the parameter v captures not only drilling outcomes but also any costs that are not directly associated with

alliance and the solo firm submit a bid of b , implying that the alliance's signal is $b_1^{-1}(b)$ while the solo firm's signal is $b_2^{-1}(b)$. If the alliance develops the lease, it receives a payoff of $fb_1^{-1}(b) + (1-f)b_2^{-1}(b) + v - b$; if the solo firm develops the lease, it receives a payoff of $fb_1^{-1}(b) + (1-f)b_2^{-1}(b) - b$. Intuitively, comparing outcomes when the two bidders submit the same bid controls for the bidders' signals and therefore controls for lease type. The only difference between the outcomes for the two bidders is then driven by their differences in productivity. While the alliance's ex ante informational advantage regarding lease type affected equilibrium bidding strategies and therefore affected the range of lease types for which the two bidders submit the same bid, conditioning on $R^* = 1$ isolates the effect of the bidders' ex post productivity differences from the effect of their differences in ex ante information.

The model presented here makes several simplifying assumptions to obtain closed-form solutions, but the basic insights that it illustrates regarding the regression discontinuity strategy generalize to other settings. Changing the signal distributions and the functional form for the common component of the lease's value would alter equilibrium bidding strategies, but conditioning on $R^* = 1$ would still control for the bidders' signals and thus would still control for lease heterogeneity. Introducing more than one alliance bidder and more than one solo firm bidder in the auction would also alter bidding strategies, but again, conditioning on the highest alliance bidder and the highest solo firm bidder submitting the same bid would still control for lease heterogeneity, as the procedure holds fixed the signals of the two highest bidders and holds fixed the distributions of the signals of the remaining bidders.

drilling, such as the cost of forming and maintaining an alliance. The variables *MEANBHNPV* and *LSNPV* should therefore be viewed as measuring an important component of the quantity v .

The model also helps to clarify the interpretation of the estimates produced by the empirical strategy. As is the case with the results of any regression discontinuity design, the estimates of the differences in drilling outcomes between solo firms and alliances are “local” to the discontinuity—that is, the estimates apply to leases for which the highest solo firm bid and the highest alliance bid are the same. If the difference in drilling outcomes between solo firms and alliances is not fixed but instead varies with lease type, the regression discontinuity strategy estimates an average treatment effect among leases for which $R^* = 1$. Without further assumptions, it is difficult to say much more about the lease types that play prominent roles in this average treatment effect. However, it seems reasonable to believe that low weight is given to leases where solo firms have important net advantages or disadvantages relative to alliances, since solo firms and alliances are unlikely to submit close bids in these cases. Conversely, high weight is probably given to leases where solo firms and alliances have comparable prospects for success.

In addition, the model highlights the key assumptions that are necessary for the regression discontinuity strategy to be valid. In order to estimate the difference in outcomes at the discontinuity $R^* = 1$, it is necessary to use data in a neighborhood of the discontinuity and to assume that the outcomes for winning solo firms and the outcomes for winning alliances are continuous functions of R^* as R^* approaches one. It is impossible to test directly whether these assumptions are valid, but I follow the guidance offered by Imbens and Lemieux (2008) and provide indirect evidence that supports the assumptions. There are two broad sets of reasons why the assumptions may not hold.

First, it would be problematic if the auction participants manipulate their bidding strategies so that R^* is just barely on one side of the discontinuity or the other. For

example, if solo firms and alliances colluded such that solo firms always won the auctions for certain lease types and alliances always won the auctions for other lease types, with non-winning parties submitting slightly lower bids in order to give the appearance of a competitive auction, the outcomes for winning solo firms and the outcomes for winning alliances would not be continuous in R^* as R^* approaches one, and the regression discontinuity strategy would not be valid. Fortunately, it does not seem likely that such collusion took place in Gulf of Mexico lease auctions. Presumably, a collusive arrangement would have been designed to allow the colluders to win auctions with low bids, with the illusion of competition intended to deter the Minerals Management Service from rejecting the winning bids as too low. This arrangement seems difficult to sustain. The colluding parties as well as outside bidders would have financial incentives to offer slightly higher bids and win the auctions. Furthermore, previous work that has tested for collusion in Gulf of Mexico lease auctions has not found evidence of it (Hendricks, Porter, and Boudreau, 1987; Hendricks, Porter, and Tan, 2003). It is also possible to examine the bidding data for indirect evidence of this type of manipulation. Figs. 1 and 2 show the distribution of leases according to their values of R^* . These figures do not reveal the presence of systematic manipulation, as the leases close to the discontinuity appear to be evenly distributed on either side of the discontinuity.

Second, it would be problematic if observable lease characteristics change discontinuously at $R^* = 1$, as such a pattern would suggest that the regression discontinuity strategy is failing to control for lease heterogeneity. For the regression discontinuity design to be valid, the explanatory variable of interest—in this case, the

assignment of a lease to a solo firm or to an alliance—must change discretely at the discontinuity while other covariates having to do with lease type change continuously. In the context of Gulf of Mexico lease auctions, it is clear that assignment of a lease to a solo firm or an alliance, which is by its nature a discrete variable, changes exactly at $R^* = 1$. I inspect several lease-level variables related to lease type for discrete changes at $R^* = 1$ and do not find evidence that they change discontinuously.

For example, I estimate using a local linear regression and the regression bandwidth proposed by Imbens and Kalyanaraman (2012)¹⁸ that the winning auction bid per acre jumps by \$866 with a standard error (s.e.) of \$747 at $R^* = 1$. This estimate is not statistically significant and is small relative to the variable's overall standard deviation (s.d.) in the data set of \$5,230. The jump at $R^* = 1$ is similarly small for the number of active leases in the same MMS block or in an adjacent MMS block at the time of the auction (estimate -0.06, standard error 0.22, overall standard deviation of variable 1.14), the number of active leases that had produced oil or gas in the same MMS block or in an adjacent MMS block at the time of the auction (estimate 0.06, s.e. 0.14, overall s.d. 0.66), the minimum water depth of the MMS block (estimate 1 foot, s.e. 20 feet, overall s.d. 109 feet), the maximum water depth of the MMS block (estimate -53 feet, s.e. 42 feet, overall s.d. 162 feet), the number of solo firm bidders (estimate -0.23, s.e. 0.35, overall s.d. 1.91), the number of alliance bidders (estimate 0.10, s.e. 0.29, overall s.d. 2.11), the amount of money submitted as bids by solo firms (estimate \$526 per acre, s.e. \$1,511 per acre, overall s.d. \$6,302 per acre), and the amount of money submitted as bids by alliances (estimate \$3,599 per acre, s.e. \$2,314 per acre, overall s.d. \$16,095 per acre).

¹⁸ See Section 4.2 for a description of this method.

Graphical analysis corroborates the statistical analysis. Figs. 3 and 4 plot the mean winning bid for leases that fall in a given range of R^* . The evidence suggests that the winning auction bid changes continuously across the discontinuity. The conclusion is similar for the other variables related to lease type. Thus, the evidence indicates that covariates related to lease type do not change discretely at $R^* = 1$, supporting the argument that the regression discontinuity strategy successfully controls for lease heterogeneity.

4.2. Implementation of the regression discontinuity strategy

To implement the regression discontinuity strategy described in the previous subsection, I use local linear methods, as suggested by Imbens and Lemieux (2008). The regression model is:

$$OUTCOME_i = \alpha + \beta \cdot (1 - BIDRATIO_i) + \gamma \cdot SOLO_i + \delta \cdot (1 - BIDRATIO_i) \cdot SOLO_i + \zeta \cdot X_i + \varepsilon_i. \quad (8)$$

Here, i indexes leases; $OUTCOME_i$ is an outcome measure for lease i ; $BIDRATIO_i$ is as defined previously; $SOLO_i$ is an indicator variable that takes a value of one if lease i was acquired by a solo firm and a value of zero if it was acquired by an alliance; X_i is a vector of control variables; and ε_i is an error term. The equation is estimated using ordinary least-squares regressions that restrict the sample to observations for which $BIDRATIO_i$ is within a bandwidth h of the discontinuity (that is, $1 - BIDRATIO_i < h$). The regressions fit a linear function of $BIDRATIO_i$ on either side of the discontinuity, and the coefficient γ on $SOLO_i$ is the parameter of interest, as it represents the difference in outcomes between solo firms and alliances at the discontinuity.

The main results in this paper implement the regression discontinuity strategy using the bandwidth h shown to be optimal for minimizing an expected squared error loss criterion by Imbens and Kalyanaraman (2012), who provide code for calculating their proposed bandwidth. As advocated by Imbens and Lemieux (2008), I also check the robustness of the main results by performing local linear regressions that use twice the suggested bandwidth and half the suggested bandwidth.

4.3. Main results: borehole profitability and lease profitability

I begin my analysis of investment outcomes by examining *MEANBHNPV*, which is the mean of borehole profitability for all boreholes drilled on a given lease within five years of the lease start date. Note that drilling did not occur on 146 out of the 1,070 leases in the primary sample. These 146 leases are excluded from the analysis of *MEANBHNPV*. Fig. 5 gives an overall sense of how *MEANBHNPV* varies with *BIDRATIO* and the status of the winning bidder as a solo firm or alliance. Fitted values from two local linear regressions, one using the Imbens-Kalyanaraman bandwidth and the other using twice that bandwidth, are superimposed on the plot of means of *MEANBHNPV* by *BIDRATIO* bin. *MEANBHNPV* generally appears to be lower for leases where solo firms won the bidding than for leases where alliances won the bidding.

Table 4 displays local linear regression results that quantify these differences in outcomes between solo firms and alliances. The middle two columns use the Imbens-Kalyanaraman bandwidth; the first two columns use twice that bandwidth; and the last two columns use half that bandwidth. The first regression within each pair includes no control variables, and the second adds the logarithm of the winning bid per acre as well as lease start year fixed effects as controls. The coefficient estimates indicate that

MEANBHNPV is approximately \$1.6 million lower (in 1980 dollars) for solo firms than for alliances at the discontinuity. This difference is economically meaningful—approximately half of the standard deviation of *MEANBHNPV*—as well as statistically significant, although imprecisely estimated. The point estimates for the effect become larger in magnitude as the bandwidth shrinks, but for a given bandwidth the estimates are largely unaltered by the inclusion of control variables. Conditional on proceeding with drilling on a lease, there is evidence that alliances have superior drilling outcomes compared to solo firms.

The results for profitability at the lease level mirror the results for mean borehole profitability. The analysis no longer conditions on drilling having occurred on a lease, but instead tries to measure the overall profitability of a lease (including the cost of the winning bid). Fig. 6 is the same as Fig. 5 except that it has lease net present value *LSNPV* as the outcome variable. Similarly, Table 5 parallels Table 4 but uses *LSNPV* as the left-hand-side variable in the local linear regressions. For Fig. 6 and Table 5, the Imbens-Kalyanaraman bandwidth is recalculated using the relevant outcome measure.

The regression discontinuity estimates in Table 5 indicate that leases where a solo firm won the auction are less profitable than leases where an alliance won the auction by \$31 million (in 1980 dollars). This effect is approximately half of the standard deviation of *LSNPV*. As with the results for mean borehole profitability, the estimates for lease profitability are statistically significant and are larger in magnitude when the bandwidth is smaller, but the addition of control variables does not change the effect size when the bandwidth is held fixed.

To further investigate the sensitivity of the results to the choice of bandwidth, I perform simple t -tests that compare the mean outcome when a solo firm wins the auction to the mean outcome when an alliance wins the auction, restricting the sample to leases for which *BIDRATIO* is greater than a progressively higher threshold value. Using *MEANBHNPV* as the outcome variable and restricting the sample to leases for which *BIDRATIO* is greater than 0.9, the mean outcome for solo firms is \$2.3 million lower (in 1980 dollars) than the mean outcome for alliances, and the p -value for the statistical test of the difference is 0.01. With a *BIDRATIO* threshold of 0.95, the difference is \$1.6 million (p -value 0.04), and with a *BIDRATIO* threshold of 0.975, the difference is \$2.0 million (p -value 0.16). When the outcome variable is *LSNPV* and the *BIDRATIO* threshold is 0.9, the mean outcome for solo firms is \$42 million lower than the mean outcome for alliances (p -value 0.01). Using a *BIDRATIO* threshold of 0.95 gives a difference of \$35 million (p -value 0.05), and using a *BIDRATIO* threshold of 0.975 gives a difference of \$57 million (p -value 0.09). Thus, for both outcome variables, increasing the *BIDRATIO* threshold reduces statistical power, but the estimated differences between solo firms and alliances deliver a qualitatively similar message.

5. Possible sources of the performance advantage of alliances

5.1. Alliances' ability to combine the information and expertise of their members

The main results presented in Section 4 provide evidence that alliances have better oil and gas drilling outcomes than solo firms. To better understand these main results, I decompose the outcome variables *MEANBHNPV* and *LSNPV* into operating profits and drilling costs (the component of *LSNPV* driven by the amount of the winning

bid is ignored in this analysis). The operating profit component of *MEANBHNPV* is the mean of borehole operating profits *BHOP*, defined in Section 3.1, taken over all boreholes drilled on a lease within five years of the lease start date. The drilling cost component of *MEANBHNPV* is the mean of borehole drilling and equipping costs *BHCOST*, also defined in Section 3.1, taken over the same set of boreholes. For *LSNPV*, the operating profit component is the discounted sum of *BHOP* for all boreholes drilled on a lease within 20 years of the lease start date, and the drilling cost component is the discounted sum of *BHCOST* for those same boreholes.

Table 6 reports the results of local linear regressions with the operating profit and drilling cost components of *MEANBHNPV* and *LSNPV* as outcome variables (columns 2, 3, 5, and 6). For comparability, the regressions use the Imbens-Kalyanaraman bandwidths from Tables 4 and 5. The estimated differences in outcomes between solo firms and alliances reveal that the results in Tables 4 and 5 are driven by alliances' differential success at generating operating profits, as operating profits are lower for solo firms than for alliances at the discontinuity.¹⁹ The negative estimates for the difference between solo firms and alliances in the drilling cost columns indicate that solo firms spend less on drilling than alliances at the discontinuity, an effect pushing in the opposite direction of the main results.

In columns 1 and 4 of Table 6, I further break down operating profits and focus on the extent to which differences in operating profits are driven by the quantity of oil and gas produced, as opposed to changes in oil and gas prices. I repeat the calculations of the operating profit components of *MEANBHNPV* and *LSNPV*, except I ignore

¹⁹ This result does not seem to be an artifact of my decision to measure oil and gas production only over the first 25 years of a borehole's life. The boreholes drilled by alliances are more likely to be productive after 25 years than those drilled by solo firms.

operating costs and royalty payments and convert production into monetary terms by applying oil and gas prices from 1960 (measured in 1980 dollars) instead of applying oil and gas prices prevailing at the time of production. Local linear regressions using these constant-price outcome measures indicate that quantities of oil and gas produced are an important driver of the performance differences between solo firms and alliances. The reported coefficients are smaller than in previous tables largely because real oil and gas prices were lower in 1960 than in other years in the sample.

Finally, to explore whether the superior performance of alliances is partly the result of alliance members pooling their information and expertise, I construct a proxy for the usefulness of such pooling, and I investigate whether the superior performance of alliances is driven by cases where the proxy indicates that pooling is most useful. Recall that the Minerals Management Service (MMS) has divided Gulf of Mexico leases into several geographic “areas” (see Section 3.1). For each set of leases that are in the same MMS area and that were auctioned simultaneously, I examine the firms that submitted bids in those auctions. For every such firm, I count the number of leases previously owned by the firm in the same MMS area, and I define a “high-experience” firm as a firm with more previous ownership experiences than the median firm that participated in those auctions. Then, for each solo firm or alliance bidder, I count the number of high-experience firms that constituted the bidder. A solo firm can be composed of zero or one high-experience firm, and an alliance can be composed of zero, one, or more high-experience firms. Table 7 reports the distribution of experience among winning solo firms and alliances for various *BIDRATIO* bandwidths. Across the bandwidths, about 40% of winning solo firms are high-experience firms. A little less than 20% of winning

alliances are composed of exactly one high-experience firm, and approximately one-third of winning alliances are composed of two or more high-experience firms. If the pooling of information and expertise plays a role in the superior performance of alliances relative to solo firms, it is likely that alliances composed of zero or one high-experience firm are more similar to solo firms in their drilling outcomes than are alliances composed of two or more high-experience firms. Below, I investigate this hypothesis.

Table 8 presents regression discontinuity estimates of performance differences between solo firms and alliances when various experience-based restrictions are placed on the data set. The first column, which has mean borehole profitability *MEANBHNPV* as the outcome variable, compares low-experience solo firms to alliances composed of zero high-experience firms. The estimation strategy is the same as the estimation strategy for the main results, except *BIDRATIO* is recalculated using only bids submitted by low-experience solo firms and alliances composed of zero high-experience firms. Thus, the regression discontinuity estimate is based on a comparison of leases where low-experience solo firms and alliances with zero high-experience members narrowly outbid each other. The regression, like all regressions in the table, includes separate linear functions of *BIDRATIO* on both sides of the discontinuity, the logarithm of the winning bid per acre in 1980 dollars, and lease start year fixed effects. For the sake of comparability, the bandwidth is the same as in the middle pair of columns in Table 4, which use the Imbens-Kalyanaraman bandwidth for the regression that does not restrict the sample on the basis of experience. The resulting estimate for the difference in *MEANBHNPV* between low-experience solo firms and alliances composed of zero high-

experience firms is +0.1 million 1980 dollars, compared to -1.6 million 1980 dollars when the sample is not restricted on the basis of experience.

The second column of Table 8 repeats this procedure but compares high-experience solo firms to alliances composed of exactly one high-experience firm. The estimated difference in *MEANBHNPV* at the discontinuity is again +0.1 million 1980 dollars, but the standard error is large because of the small sample size. To improve power, the third column pools alliances composed of zero or one high-experience firm and compares this group to the group of all solo firms (both low-experience and high-experience). The estimated difference is -0.7 million 1980 dollars, a 60% decrease in magnitude relative to the difference of -1.6 million 1980 dollars when there are no experience-based restrictions placed on the sample. To assess the statistical significance of this decrease, I use the following bootstrap procedure. I create ten thousand bootstrap samples by randomly sampling the data set with replacement. For each bootstrap sample, I obtain regression discontinuity estimates both with and without experience-based restrictions placed on the sample, and I calculate the difference between these estimates. The distribution of the difference across bootstrap samples indicates that the difference is statistically significant at the 5% level.

In the fourth column of Table 8, I use the methodology above to compare the group of all solo firms (both low-experience and high-experience) to the group of alliances composed of two or more high-experience firms. The regression discontinuity estimate of the difference in *MEANBHNPV* is -3.4 million 1980 dollars, more than twice the estimate from the regression that does not place experience-based restrictions on the sample. This difference in estimates is statistically significant at the 10% level, and the

difference between the estimates in the third and fourth columns of Table 8 is statistically significant at the 5% level.

The remaining columns of Table 8 perform the same analysis using lease net present value *LSNPV* instead of mean borehole net present value *MEANBHNPV* as the outcome variable. The results are similar, although none of the differences in regression discontinuity estimates across sample definitions are statistically significant.

Overall, the evidence suggests that the differences in drilling outcomes between solo firms and alliances composed of zero or one high-experience firm are smaller than the differences between solo firms and alliances composed of two or more high-experience firms. These findings are consistent with the hypothesis that the superior performance of alliances is due to alliances giving their member firms the opportunity to combine information and expertise. According to this hypothesis, when alliance members can each contribute insights into the decision of where to drill within a lease, the drilling location chosen will produce more oil and gas than the location that would have been chosen by a solo firm. Alliances then take advantage of this ability to achieve higher borehole profitability by engaging in more drilling on average than solo firms. Finally, alliances' higher borehole profitability and higher drilling expenditures together generate higher lease profitability relative to solo firms. Of course, while the results are consistent with this hypothesis, there are other possible explanations for the performance differences between alliances and solo firms. I now turn to a discussion of some of these possibilities.

5.2. Agency problems resolved by alliances

It is possible that alliances overcome agency problems that lead to suboptimal drilling decisions among solo firms, but the simplest versions of agency considerations are not consistent with the data. The evidence in Table 6 that solo firms have slightly lower drilling expenditures than alliances suggests that alliances do not overcome a tendency for solo firm managers to spend more on drilling than is optimal. At the same time, logit regressions modeling the probability that any drilling occurred on a lease within five years of the lease start date indicate that solo firms are, if anything, more likely than alliances to proceed with initial drilling at the discontinuity.²⁰ This finding casts doubt on the hypothesis that alliances help managers overcome excessive conservatism in their drilling decisions that is driven by, for example, a fear that a good decision with a bad outcome ex post will be harmful to their career prospects (Holmstrom, 1999).

Even if alliances do not lead managers to undertake projects that they might have otherwise forgone, the risk-sharing provided by alliances may encourage managers to develop leases using strategies that have higher risk but higher expected returns (Palia, Ravid, and Reisel, 2008). To evaluate this hypothesis, I perform a series of quantile regressions that estimate the differences at the discontinuity between solo firms and alliances for various points in the distribution of lease profitability $LSNPV$. For example, the 0.8 quantile regression estimates the difference between the 0.8 quantile of solo firm outcomes at the discontinuity and the 0.8 quantile of alliance outcomes at the

²⁰ These results are available from the author upon request.

discontinuity. The regressions are implemented using local linear methods.²¹ I am not aware of a bandwidth selection algorithm designed for this setting, so I use the Imbens-Kalyanaraman bandwidth from Table 5.

The regression results are presented graphically in Fig. 7. Point estimates for the difference between solo firm outcomes and alliance outcomes are shown for the 0.2 quantile through the 0.8 quantile, with 95% confidence intervals depicted as dotted lines. More extreme quantiles are omitted because the confidence intervals become quite large.

There is little evidence that the lease development strategies pursued by alliances are riskier than those pursued by solo firms. The estimates for the differences between solo firm outcomes and alliance outcomes, although often not significantly different from zero, are uniformly negative, even for lower quantiles of the distribution than those shown in the figure. Furthermore, the superior performance of alliances relative to solo firms does not seem to be driven only by the extreme right tail of outcomes, since the differences between solo firm outcomes and alliance outcomes are significantly different from zero starting at approximately the 0.6 quantile.

5.3. Bidding strategies of alliances and solo firms

Another factor that may be responsible for alliances' better drilling outcomes is the possibility that alliances are less aggressive than solo firms in their auction bidding strategies. If alliances and solo firms have similar drilling profitability but there is a greater systematic difference between an alliance's bid and its drilling profitability than between a solo firm's bid and its drilling profitability, leases for which the highest

²¹ For an analysis of local linear methods in quantile regression discontinuity designs, see Frandsen, Frölich, and Melly (2012).

alliance bid and the highest solo bid are nearly identical will be associated with higher drilling profitability when the alliance wins the auction.

Section 4.1 emphasized that the regression discontinuity strategy identifies a treatment effect that is local to a subset of lease types, namely, those that are likely to be present when the highest alliance bid and the highest solo bid are similar (that is, when *BIDRATIO* is close to one). Therefore, to examine the possibility that bidding strategies explain the superior drilling outcomes of alliances at the discontinuity, it is necessary to study bidding behavior in auctions for these lease types. The most direct way to pinpoint these lease types in the data is to look at leases for which *BIDRATIO* is close to one. However, since these leases are chosen precisely because their bids exhibit a special property, analyzing bids for only these leases gives an incomplete characterization of bidding behavior for the relevant lease types. To address this difficulty while maintaining a narrow focus on the relevant lease types, I examine leases for which *BIDRATIO* is greater than 0.9, together with leases that are nearby geographic neighbors of leases for which *BIDRATIO* is greater than 0.9. Recall that the Minerals Management Service has divided the Gulf of Mexico into geographic “blocks” of approximately 5,000 acres each (see Section 3.1). Two leases are nearby geographic neighbors if they are located in the same block or if they are located in different blocks with boundaries that touch. Using this definition for neighbors maintains a narrow focus on the relevant lease types because geographic proximity is associated with a high degree of similarity in lease characteristics.

There are 417 leases that have *BIDRATIO* greater than 0.9 or that are nearby geographic neighbors of leases with *BIDRATIO* greater than 0.9. In the auctions for these

417 leases, the mean solo firm bid is \$1,319 per acre (in 1980 dollars), and the mean alliance bid is \$2,634 per acre (in 1980 dollars), a difference that is statistically significant at the 1% level in a two-sample t -test. Fig. 8 displays empirical cumulative distribution functions of alliance bids and solo firm bids for this set of leases, and it is apparent that the two distributions have the empirical analogue of a first-order stochastic dominance relationship. This evidence indicates that alliances have more aggressive bidding strategies than solo firms for the relevant set of lease types, suggesting that differences in bidding strategies cannot explain the superior drilling performance of alliances at the discontinuity.

Because the highest alliance bid and the highest solo firm bid in an auction are the relevant bids for the purposes of the regression discontinuity strategy, I repeat the analysis for the 417 leases identified above but restrict attention to these bids. If no alliance submits a bid in a given auction, the highest alliance bid in that auction is deemed to be zero; similarly, if no solo firm submits a bid in a given auction, the highest solo firm bid in that auction is deemed to be zero. The mean highest solo firm bid is \$1,882 per acre (in 1980 dollars), and the mean highest alliance bid is \$2,329 per acre (in 1980 dollars). In a two-sample t -test, this difference is statistically significant at the 10% level. Fig. 9 shows empirical cumulative distribution functions of highest alliance bids and highest solo firm bids. Alliances are less aggressive than solo firms in the sense that alliances are less likely to submit a bid at all. However, this particular form of greater conservatism on the part of alliances cannot explain the regression discontinuity results, since the regression discontinuity estimates are based on leases for which at least one alliance and at least one solo firm submit bids. Beyond alliances' and solo firms'

different propensities to participate in auctions, the overall bid distributions of alliances and solo firms in Fig. 9 are quite similar, but alliance bids are generally slightly more aggressive than solo firm bids. Again, the evidence suggests that the main regression discontinuity results are not due to differences in alliance and solo firm bidding strategies.

5.4. Selection of firms into alliances

Another potential explanation for the main results is that high-quality firms are more likely to form alliances, perhaps because it is easier for them to find partners that are willing to work with them. To investigate this possibility, I introduce firm fixed effects into the regression analysis. Recall from Table 1 that many companies participated in lease auctions both as solo firms and as alliance members, so it is possible to compare companies working as solo firms to the same companies working as part of alliances.

I construct two different sets of firm effects. The first set, which I label “lead firm indicators,” is a set of dummy variables, one for each company, that take a value of one if the company in question was the lead firm for the winning lease bid. In cases where a solo firm won the auction, the lead firm is simply the winning bidder. In cases where an alliance won the auction, the lead firm is defined as the firm with the largest percentage ownership share of the winning bid. Ties are broken by assigning lead firm status to the company that contributed the largest number of 1980 dollars to winning bids in the sample.

I label the second set of firm effects “percent ownership variables.” Each company has a distinct percent ownership variable, which takes a value equal to the company’s percent ownership of the winning lease bid.

Table 9 reports the results of regressions that incorporate these firm effects. The outcome variable is mean borehole profitability *MEANBHNPV* or lease profitability *LSNPV*. The regressions use a bandwidth that is twice the one suggested by Imbens and Kalyanaraman (2012). The Imbens-Kalyanaraman algorithm is designed to provide the optimal bandwidth for regression discontinuity implementations with a small number of control variables. I employ a wider bandwidth in the current setting so that the firm effects can be estimated more reliably. The estimates of profitability differences between solo firms and alliances at the discontinuity are not altered substantially by the inclusion of the firm effects, and in some cases they are larger in magnitude.

To account for the possibility that firm quality changes over time during the sample period, Table 9 also displays the results of regressions that interact the firm effects with dummies for the five-year window containing the lease start date.²² The finding that alliances have better drilling outcomes than solo firms persists, although the results are slightly weaker.

Finally, although it is impossible to assess the effect of all potentially relevant firm attributes on the propensity to select into alliances, I can provide some evidence on the nature of the selection process. The results presented in Section 5.1 suggest that a firm’s past experience in the vicinity of a lease can help the firm achieve better drilling outcomes on the lease, so it is informative to explore how past experience influences a

²² I have experimented with interacting the firm effects with dummies for one-year lease start date windows, but the results are generally uninformative because of the large standard errors on the coefficient estimates.

firm's decision of whether to join an alliance. For this analysis, I focus on leases that are relevant for the regression discontinuity estimates by limiting attention to leases with *BIDRATIO* greater than 0.9. Using the definition of high-experience firms from Section 5.1,²³ I find that when a high-experience firm participates in one of these lease auctions, it submits a solo firm bid in 32% of the cases. In contrast, when a low-experience firm participates in one of these lease auctions, it submits a solo firm bid in 23% of the cases. Thus, high-experience firms submit solo firm bids more frequently than low-experience firms, a pattern which is consistent with the argument that a high-experience firm may be more likely than a low-experience firm to prefer earning all of the profits from developing a lease on its own as opposed to earning a fraction of the (potentially larger) profits from developing a lease as part of an alliance (Hendricks, Porter, and Tan, 2008). This selection process tends to push against the finding that alliances have better drilling performance than solo firms.

Overall, the evidence suggests that the selection of firms into organizational types is not the primary force behind the main results.

5.5. Differences in the discount rates of alliances and solo firms

Differences in the discount rates applied to drilling projects by alliances and solo firms, perhaps induced by differences in financial constraints, could account for the main performance results. More generally, the main results could reflect differences in the financial resources available to alliances and solo firms. To explore these possible explanations, I match firms in the sample to Compustat in order to calculate variables that

²³ Recall that the definition applies at the firm-lease level. A given firm may be a high-experience firm when bidding in one lease auction but a low-experience firm when bidding in another lease auction.

capture the leverage, profitability, and available resources of winning auction bidders, and I include these variables in the regression analysis as control variables.

Market leverage for a winning solo firm bidder is defined as the book value of the firm's debt (the sum of long-term debt and debt in current liabilities), divided by the sum of the book value of the firm's debt and the market value of the firm's equity, calculated as of the end of the fiscal year immediately preceding the lease auction. Market leverage for a winning alliance bidder is defined as the weighted average of market leverage for its constituent firms, where the market leverage for a constituent firm is calculated as above and where the weight is given by the firm's percent ownership of the winning bid. When a constituent firm could not be matched to Compustat, it is excluded from the calculation, and the weights of the remaining firms are increased proportionately.

The profitability of a solo firm is calculated as the most recent fiscal year's operating income after depreciation, divided by the firm's total assets as of the end of the most recent fiscal year. Alliance profitability is the weighted average of the profitability of constituent firms, applying the same weighting conventions used to calculate alliance market leverage.

The available cash of a solo firm is defined as the amount of cash and short-term investments, in 1980 dollars, on the firm's balance sheet at the end of the most recent fiscal year. The total assets of a solo firm are also observed at the end of the most recent fiscal year and converted to 1980 dollars. For an alliance, the available cash is the sum of the available cash of constituent firms, and the total assets are the sum of the assets of constituent firms, in both cases disregarding firms that could not be matched to Compustat.

The first column of Table 10 reports the results from a regression that has the same specification and the same bandwidth as the regression reported in the fourth column of Table 4, which presents the main results with mean borehole profitability *MEANBHNPV* as the outcome variable. However, the regression sample in Table 10 does not include observations for which market leverage, profitability, available cash, or total assets is missing due to a lack of a match in Compustat. The third column of Table 10 adopts the same sample restriction as the first column of Table 10 but otherwise has the same regression specification and bandwidth as the fourth column of Table 5, which presents the main results with lease profitability *LSNPV* as the outcome variable. The sample restriction makes the regression discontinuity estimates in Table 10 slightly larger in magnitude than but qualitatively similar to the regression discontinuity estimates in Tables 4 and 5.

The regressions reported in the second and fourth columns of Table 10 add market leverage, profitability, the logarithm of available cash, and the logarithm of total assets, all winsorized at their respective 1st and 99th percentiles, as control variables. Although the coefficient estimates on these control variables are not statistically significant, they are generally of the predicted sign. Consistent with the argument that higher leverage is associated with higher discount rates and therefore higher profitability as measured by the outcome variables *MEANBHNPV* and *LSNPV*, the coefficient estimate on the market leverage variable is positive. The coefficient estimates on the profitability and available cash variables are also positive, in line with the hypothesis that financial resources improve drilling outcomes. The coefficient estimate on the total assets variable is negative, although this result is a partial correlation—controlling for the

profitability and available cash variables, an increase in total assets may reflect a larger number of projects that require financial resources. Importantly, after the introduction of these control variables, the regression discontinuity estimates are negative, statistically significant, and of approximately the same magnitude. The control variables are imperfect measures of discount rates and financial resources, but these regression results suggest that differences in discount rates and financial resources are not the primary driving force behind the main results.

5.6. Systematic measurement error in the cost of drilling

As noted in Section 3.1, the outcome variables *MEANBHNPV* and *LSNPV* are calculated using estimates of drilling costs that are subject to measurement error. To estimate the cost of drilling a given borehole, I multiply the number of feet drilled, which is measured quite precisely, by the average cost per foot of drilling similar boreholes, where similarity is based on time period, geographic location, type of well, and well depth. Thus, the actual cost of drilling a given borehole need not match the estimated cost of drilling the borehole.

To examine the possibility that the main empirical results are driven by systematic measurement error, I explore the sensitivity of the results to changes in drilling cost assumptions. If alliances are assumed to have drilling costs per foot that are 25% higher than the drilling costs of solo firms, the main regression discontinuity estimates remain similar in magnitude and statistically significant. In order to explain the estimated differences in *MEANBHNPV* and *LSNPV* between alliances and solo firms at the discontinuity, the drilling costs of alliances would have to be more than twice the drilling costs of solo firms.

It is unlikely that the drilling costs of alliances are more than 25% higher than the drilling costs of solo firms. Conditioning on the year, region, well type, and well depth of a borehole reduces the scope for differences in cost.²⁴ Furthermore, as can be seen in Table 1, many firms conduct drilling both as part of an alliance and as a solo firm, again reducing the scope for systematic differences in costs between alliances and solo firms. Thus, it is unlikely that systematic measurement error in drilling costs explains the main empirical results.

5.7. Incentives for information production

An alternative set of explanations for the main results, which is difficult to test, is the possibility that alliances do not simply combine the information and expertise of member firms but instead create stronger incentives for new information and expertise to be acquired. There are several possible mechanisms by which alliances may create stronger incentives for information production. Alliances may commit firms to provide resources to support the development of a particular lease, while solo firms retain the ability to redirect resources to other projects. The manager responsible for developing a lease on behalf of an alliance might be more motivated to invest in new information and expertise relevant to the lease with the knowledge that such efforts will not be wasted (Stein, 1997; Robinson, 2008). It is also possible that alliances improve managerial effort by providing increased monitoring. On the other hand, by giving their member firms only fractional ownership stakes in a lease, alliances may weaken member firms' incentives to exert effort to overrule the decisions of the manager responsible for

²⁴ In contrast, changing the exact location of a borehole within the boundaries of a lease can generate a large change in the revenue produced by the borehole, based on the presence or absence of oil and gas.

developing the lease, thereby increasing the manager's willingness to invest in improving the quality of those decisions (Aghion and Tirole, 1997).

I cannot observe the internal governance of alliances or solo firms, so it is difficult to provide evidence that supports or refutes this set of hypotheses. However, these hypotheses would not necessarily predict that the superior performance of alliances is related to the previous experience of member firms in the geographic vicinity of the lease, as I document in Section 5.1. Furthermore, these alternative explanations are not very different in spirit from the explanation that alliances combine the information and expertise of their member firms to improve drilling outcomes. In both cases, alliances serve as a mechanism for bringing more information and expertise to bear on an investment project.

6. Conclusion

Using a regression discontinuity strategy to control for lease heterogeneity, I find that oil and gas firm alliances in the Gulf of Mexico have higher mean borehole profitability and higher lease profitability than solo firms. The results are consistent with the hypothesis that alliances achieve superior outcomes by combining the information and expertise of their members.

These findings support theoretical arguments in the incomplete contracts literature on the boundaries of the firm that alliances are an effective organizational design for fine-tuning the incentives of member firms to contribute to joint investment projects (Robinson, 2008; Fulghieri and Sevilir, 2009; Hackbarth, Mathews, and Robinson, 2012). While the setting for this paper's empirical analysis, oil and gas drilling in the Gulf of

Mexico, is a special environment, the insights regarding the function of alliances in facilitating the sharing of information and expertise are likely to generalize to other contexts, such as the biotechnology and telecommunications industries, in which noncontractible inputs play an important role in joint production efforts. Furthermore, the motivations for alliance creation that are considered here are likely to be at play in other cases of changing firm boundaries, such as mergers, acquisitions, the formation of conglomerates, divestitures, and spinoffs, all of which may reflect shifts in the opportunities for distinct organizational units to pursue investment projects that require their combined capabilities. Future research should recognize the sharing of information and expertise as a key determinant of organizational form and further explore its role in driving financial outcomes.

Appendix. Calculation of economic returns from drilling

I use a modified version of the approach of Hendricks, Porter, and Boudreau (1987) to calculate economic returns from drilling activity.

I start with borehole-level data on the total borehole vertical depth, the distance along the borehole's axis from the drilling rig to the bottom, the water depth at the location, the latitude and longitude of the borehole bottom, the latitude and longitude of the drilling rig, the elevation of the drilling rig above water level, and the spud date (the date on which drilling commenced). There are two main complications that arise when converting this information into the total feet of penetration down the borehole. First, a borehole is not necessarily straight but can bend to avoid drilling through problematic geological formations. I disregard this possibility and idealize each borehole as a straight line from one point to another. Second, a borehole may be a "sidetrack" of a previously drilled borehole. A sidetrack is created by drilling through the side wall of an existing borehole and proceeding to a new bottom location. The initial portion of a sidetrack coincides with a segment of a pre-existing borehole, so a sidetrack requires less drilling than might otherwise be suggested by its length. There are 411 (1,145) sidetracks in my sample of 5,833 (10,415) boreholes drilled within five (20) years of the lease start date. My methodology for working with sidetracks is explained below.

I arrange boreholes into groups such that each group contains one initial borehole and all of its sidetracks (and any sidetracks of sidetracks). I assume that the initial borehole is drilled on the straight line extending from the drilling rig to the borehole bottom, and I assume that the point where this line intersects the sea floor is the top of the borehole. I calculate the total feet of penetration down this borehole as the distance along

the borehole's axis from the drilling rig to the bottom, minus the distance from the drilling rig to the top of the borehole. Note that my calculation of the total feet of penetration may exceed the straight-line distance between the assumed borehole top and the borehole bottom. This discrepancy may arise even when the assumed borehole top matches the true borehole top (which is not observed in the data), since the borehole may bend on its way to the bottom instead of proceeding in a straight line. Also, depending on the location of the assumed borehole top in relation to the true borehole top, my calculation may overstate or understate the true total feet of penetration. However, this imprecision is likely to be minimal because most drilling during this time period occurred along a straight line.

For all sidetracks in a group, I calculate the total feet of (incremental) penetration as follows. I take each previously drilled borehole in the group, and I consider every point on the line segment between the pre-existing borehole's top and bottom as a "candidate" for the location at which incremental drilling for the new sidetrack commences. Assuming that incremental drilling proceeds along the straight line extending from this "candidate location" to the sidetrack bottom, every "candidate location" implies a distance along the axis of the borehole from the drilling rig to the sidetrack bottom. For the "most reasonable candidate location," I choose the "candidate location" that minimizes the difference between this implied distance and the actual distance reported in the data, subject to the constraint that the "candidate location" cannot have a greater vertical depth than the sidetrack bottom. The total (incremental) penetration is then calculated as the straight-line distance between the "most reasonable candidate location" and the sidetrack bottom. In cases where there are multiple

previously drilled boreholes in the group, I choose from the set of “most reasonable candidate locations” by minimizing the incremental distance drilled.

To calculate drilling costs, I use data from the American Petroleum Institute (API) drilling cost survey, which has been conducted annually since the 1950s. Survey data on the cost per foot drilled are separately reported by region (offshore Louisiana or offshore Texas), well type (oil, gas, or unproductive), and well depth.²⁵ When cost per foot data are unavailable for a particular region, well type, well depth, and year, I fill in the cell using the cost per foot from the other region for the same well type, well depth, and year. When data are still missing after performing this procedure, I use the cost per foot from the nearest available well depth category for the same region, well type, and year. Finally, because the survey was not conducted in 1957 and 1958 and because I was unable to locate survey data from 1968, I interpolate assuming that the cost per foot for a given region, well type, and well depth grew at a constant exponential rate from 1956 to 1959 and at a constant exponential rate from 1967 to 1969. I classify each borehole in my data set according to its region, type, depth, and year of spud date, and I multiply the total feet of penetration by the appropriate drilling cost per foot to obtain the cost of drilling the borehole. All costs are converted to 1980 dollars using the GDP implicit price deflator from the Bureau of Economic Analysis.

In a departure from the methodology of Hendricks, Porter, and Boudreau (1987), I also adjust for the cost of lease equipment beyond the costs incorporated into the drilling calculation. This adjustment is only applied to the drilling cost of productive boreholes, since unproductive boreholes do not require additional equipment. To calculate the

²⁵ There are 11 well depth categories: less than 1,250 feet; 1,250–2,500 feet; 2,500–3,750 feet; 3,750–5,000 feet; 5,000–7,500 feet; 7,500–10,000 feet; 10,000–12,500 feet; 12,500–15,000 feet; 15,000–17,500 feet; 17,500–20,000 feet; and greater than 20,000 feet.

adjustment factor, I divide total estimated annual expenditures on lease equipment in the U.S. by total estimated annual expenditures on the drilling of productive wells in the U.S. Both of these estimates are from the API survey for the years 1955–1975. I assume that the adjustment factor grew at a constant exponential rate from 1956 to 1959 in order to interpolate for the years 1957 and 1958 (I was able to locate this component of the survey data for 1968). For the years 1976–1982, I use comparable data from the Census Bureau’s Annual Survey of Oil and Gas. Data for the years 1983–2000 are from the Energy Information Administration (EIA). However, the EIA reports lease equipment cost *indices* for oil wells and gas wells, and it does not report total expenditures on lease equipment. To convert this information into an adjustment factor comparable to the other two data series, I use the following procedure. Define Peq_{cy} as the price of a unit of lease equipment for commodity c (where c is o for oil or g for gas) in year y ; Pft_{cy} as the price per foot of drilling a productive well for commodity c in year y ; Xeq_{cy} as the number of units of lease equipment purchased for commodity c in year y ; and Xft_{cy} as the number of feet of productive wells drilled for commodity c in year y . Then the lease equipment adjustment factor LE_y , which is the ratio of total expenditure on lease equipment to total expenditure on productive wells, can be expressed as

$$\begin{aligned}
LE_y &= \frac{Peq_{oy} Xeq_{oy} + Peq_{gy} Xeq_{gy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}} \\
&= \frac{Peq_{oy} Xeq_{oy}}{Pft_{oy} Xft_{oy}} \frac{Pft_{oy} Xft_{oy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}} + \frac{Peq_{gy} Xeq_{gy}}{Pft_{gy} Xft_{gy}} \frac{Pft_{gy} Xft_{gy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}} \\
&= \frac{Peq_{oB} Xeq_{oy} \frac{Peq_{oy}}{Peq_{oB}}}{Pft_{oB} Xft_{oy} \frac{Pft_{oy}}{Pft_{oB}}} \frac{Pft_{oy} Xft_{oy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}} + \frac{Peq_{gB} Xeq_{gy} \frac{Peq_{gy}}{Peq_{gB}}}{Pft_{gB} Xft_{gy} \frac{Pft_{gy}}{Pft_{gB}}} \frac{Pft_{gy} Xft_{gy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}}.
\end{aligned}$$

Here, B is a base year, which I take to be the same across indices (this simplification is for ease of exposition and is not essential). If the number of units of lease equipment required per foot of productive well is constant across time for both commodities, then I can rewrite this equation as

$$LE_y = J_o \frac{\frac{Peq_{oy}}{Peq_{oB}}}{\frac{Pft_{oy}}{Pft_{oB}}} \frac{Pft_{oy} Xft_{oy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}} + J_g \frac{\frac{Peq_{gy}}{Peq_{gB}}}{\frac{Pft_{gy}}{Pft_{gB}}} \frac{Pft_{gy} Xft_{gy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}},$$

where J_o and J_g are constants reflecting technological parameters and the prices of lease equipment and drilling in the base year B . The fractions $\frac{Pft_{oy} Xft_{oy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}}$ and

$\frac{Pft_{gy} Xft_{gy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}}$ are simply the dollar share of productive oil well drilling and the dollar

share of productive gas well drilling, respectively, out of total expenditure on drilling

productive wells. I calculate these fractions using EIA data. The fractions $\frac{Peq_{oy}/Peq_{oB}}{Pft_{oy}/Pft_{oB}}$ and

$\frac{Peq_{gy}/Peq_{gB}}{Pft_{gy}/Pft_{gB}}$ are ratios of lease equipment price indices to drilling cost indices, which are

also available from the EIA. The EIA data overlap the Census Bureau data for the years

1976–1982, so I estimate J_o and J_g using an ordinary least squares (OLS) regression of

LE_y from the Census Bureau data on $\frac{Peq_{oy}/Peq_{oB}}{Pft_{oy}/Pft_{oB}} \frac{Pft_{oy} Xft_{oy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}}$ and $\frac{Peq_{gy}/Peq_{gB}}{Pft_{gy}/Pft_{gB}} \frac{Pft_{gy} Xft_{gy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}}$

from the EIA data (without including a constant in the regression). I then use these

estimates of J_o and J_g to extrapolate predicted values of LE_y for the years 1983–2000,

employing EIA data on $\frac{Peq_{oy}/Peq_{oB}}{Pft_{oy}/Pft_{oB}} \frac{Pft_{oy} Xft_{oy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}}$ and $\frac{Peq_{gy}/Peq_{gB}}{Pft_{gy}/Pft_{gB}} \frac{Pft_{gy} Xft_{gy}}{Pft_{oy} Xft_{oy} + Pft_{gy} Xft_{gy}}$ to generate

fitted values from the regression equation. Finally, for each productive well, I adjust the

drilling cost by multiplying it by the lease equipment cost factor $(1 + LE_y)$. These

adjusted drilling costs and the unadjusted drilling costs for unproductive wells serve as my estimates of the cost per borehole, *BHCOST*.

My procedure for calculating the cost of operating oil and gas wells is similar to my procedure for calculating lease equipment costs. For the years 1955–1975, I use data from the API survey to generate OP_y , the ratio of total estimated annual expenditures on the operation of U.S. oil and gas wells to total estimated annual revenues from oil and gas produced from U.S. wells. Comparable data from the Census Bureau allow me to calculate OP_y for the years 1976–1982. The EIA provides oil and gas operating cost indices starting in 1976, from which I construct OP_y for the years 1983–2006. Let Pop_{cy} be the price of a unit of well operations for commodity c (where c is o for oil or g for gas) in year y ; let Pv_{cy} be the price per unit volume of commodity c in year y ; let Xop_{cy} be the units of well operations used to produce commodity c in year y ; and let Xv_{cy} be the volume of commodity c produced in year y . Following logic similar to the reasoning used above to rewrite LE_y , the operating cost ratio OP_y can be expressed as

$$\begin{aligned}
 OP_y &= \frac{Pop_{oy}Xop_{oy} + Pop_{gy}Xop_{gy}}{Pv_{oy}Xv_{oy} + Pv_{gy}Xv_{gy}} \\
 &= \frac{Pop_{oy}Xop_{oy}}{Pv_{oy}Xv_{oy}} \frac{Pv_{oy}Xv_{oy}}{Pv_{oy}Xv_{oy} + Pv_{gy}Xv_{gy}} + \frac{Pop_{gy}Xop_{gy}}{Pv_{gy}Xv_{gy}} \frac{Pv_{gy}Xv_{gy}}{Pv_{oy}Xv_{oy} + Pv_{gy}Xv_{gy}} \\
 &= \frac{Pop_{oB}Xop_{oy}}{Pv_{oB}Xv_{oy}} \frac{\frac{Pop_{oy}}{Pop_{oB}}}{\frac{Pv_{oy}}{Pv_{oB}}} \frac{Pv_{oy}Xv_{oy}}{Pv_{oy}Xv_{oy} + Pv_{gy}Xv_{gy}} + \frac{Pop_{gB}Xop_{gy}}{Pv_{gB}Xv_{gy}} \frac{\frac{Pop_{gy}}{Pop_{gB}}}{\frac{Pv_{gy}}{Pv_{gB}}} \frac{Pv_{gy}Xv_{gy}}{Pv_{oy}Xv_{oy} + Pv_{gy}Xv_{gy}} \\
 &= K_o \frac{\frac{Pop_{oy}}{Pop_{oB}}}{\frac{Pv_{oy}}{Pv_{oB}}} \frac{Pv_{oy}Xv_{oy}}{Pv_{oy}Xv_{oy} + Pv_{gy}Xv_{gy}} + K_g \frac{\frac{Pop_{gy}}{Pop_{gB}}}{\frac{Pv_{gy}}{Pv_{gB}}} \frac{Pv_{gy}Xv_{gy}}{Pv_{oy}Xv_{oy} + Pv_{gy}Xv_{gy}},
 \end{aligned}$$

where B is again a base year, and where I assume that the number of units of well operations required per volume of commodity produced does not change over time for

either commodity. The expression $\frac{POP_{oy}/POP_{oB}}{Pv_{oy}/Pv_{oB}} \frac{Pv_{oy}Xv_{oy}}{Pv_{oy}Xv_{oy}+Pv_{gy}Xv_{gy}}$ is the ratio of the oil well operating cost index to the oil price index, multiplied by the dollar share of oil production out of total oil and gas production. Similarly, the expression $\frac{POP_{gy}/POP_{gB}}{Pv_{gy}/Pv_{gB}} \frac{Pv_{gy}Xv_{gy}}{Pv_{oy}Xv_{oy}+Pv_{gy}Xv_{gy}}$ is the ratio of the gas well operating cost index to the gas price index, multiplied by the dollar share of gas production out of total oil and gas production. An OLS regression of OP_y (calculated from Census Bureau data) on the two expressions above (calculated from EIA data) does not produce reliable estimates of K_o and K_g because the two expressions have a correlation of 0.82. I therefore regress OP_y on a constant and the first principal component of the two expressions. I then use the estimated regression equation, combined with EIA data, to extrapolate predicted values of OP_y for the years 1983–2006. The variable OP_y is incorporated into my calculation of the operating profits from a borehole, as explained below.

The Minerals Management Service provides data on the monthly production of oil and gas from each borehole. I multiply monthly oil production by the average offshore Louisiana wellhead crude oil price for the year, converted to 1980 dollars, and I multiply monthly gas production by the average offshore Louisiana wellhead natural gas price for the year, also converted to 1980 dollars. Wellhead prices were unavailable for the years 1955–1959, but prices were roughly stable over this period. I assume that real wellhead prices for oil and gas during this period were equal to the real wellhead prices from 1960. I then multiply monthly borehole revenue by $(1 - ROYALTY - OP_y)$, where $ROYALTY$ is the appropriate royalty rate and OP_y is the operating cost ratio calculated previously. For each borehole, I consider the stream of operating profits starting at the spud date and extending through 25 years or through the year 2006, whichever is earlier. For

subsequent years, I take the level of production from the earlier of the 25th year and the year 2006, and I assume that it declines at a 25% annual rate in perpetuity. I calculate operating profits from this production by applying oil prices, gas prices, and the operating cost ratio from the end of the 25th year or the year 2006. The variable *BHOP* is defined as the sum of all operating profit attributed to the borehole, discounted to the spud date by applying a 5% annual rate.

This calculation differs from the calculation of Hendricks, Porter, and Boudreau (1987) in three ways. First, they do not incorporate an adjustment for the operating costs of wells. Second, they multiply all production from a lease by the oil and gas prices that prevailed in the year during which the lease was auctioned. The auctions they study took place during a period of stable oil and gas prices, so it is sensible for them to examine the appropriateness of bidding strategies by assuming that bidders believed prices would remain stable. The last auctions in my sample, however, took place during a period of rapidly rising prices due to the oil embargo of 1973. Furthermore, I am primarily interested in drilling decisions subsequent to the lease auctions. My calculations therefore account for changes in prices that occur over the life of a well. Third, Hendricks, Porter, and Boudreau (1987) use data on production through 1980, and when production on a lease continued past 1980, they find the year of peak production and assume that production thereafter declined at a 25% annual rate for 15 years before dropping to zero. My alternative assumptions are simply intended to take full advantage of production data available for years beyond 1980.

My variable *BHNPV* is the present discounted value of operating profits *BHOP* minus drilling and equipping costs *BHCOST*. For a given lease, *MEANBHNPV* is the

mean of *BHNPV* among the boreholes drilled within five years of the lease start date. I also calculate total profitability at the lease level. I discount *BHNPV* back to the lease start date at a 5% annual rate, and I sum over all boreholes drilled on a lease within 20 years of the lease start date. Subtracting off the bonus paid to acquire the lease gives *LSNPV*, the net present value of the lease.

References

- Aghion, P., Tirole, J., 1997. Formal and real authority in organizations. *Journal of Political Economy* 105, 1–29.
- Anand, B.N., Khanna, T., 2000. Do firms learn to create value? The case of alliances. *Strategic Management Journal* 21, 295–315.
- Avery, C., Kagel, J.H., 1997. Second-price auctions with asymmetric payoffs: an experimental investigation. *Journal of Economics & Management Strategy* 6, 573–603.
- Baker, G.P., Gibbons, R., Murphy, K.J., 2008. Strategic alliances: bridges between ‘islands of conscious power.’ *Journal of the Japanese and International Economies* 22, 146–163.
- Baker, G.P., Hubbard, T.N., 2003. Make versus buy in trucking: asset ownership, job design, and information. *American Economic Review* 93, 551–572.
- Baker, G.P., Hubbard, T.N., 2004. Contractibility and asset ownership: on-board computers and governance in U.S. trucking. *Quarterly Journal of Economics* 119, 1443–1479.
- Berger, A.N., Miller, N.H., Petersen, M.A., Rajan, R.G., Stein, J.C., 2005. Does function follow organizational form? Evidence from the lending practices of large and small banks. *Journal of Financial Economics* 76, 237–269.
- Bertrand, M., Mullainathan, S., 2005. Profitable investments or dissipated cash? Evidence on the investment-cash flow relationship from oil and gas lease bidding. NBER Working Paper No. 11126.
- Bulow, J., Klemperer, P., 2002. Prices and the winner’s curse. *RAND Journal of Economics* 33, 1–21.
- Chan, S.H., Kensinger, J.W., Keown, A.J., Martin, J.D., 1997. Do strategic alliances create value? *Journal of Financial Economics* 46, 199–221.
- Coase, R.H., 1937. The nature of the firm. *Economica* 4, 386–405.
- Desai, M.A., Foley, C.F., Hines, J.R. Jr., 2004. The costs of shared ownership: evidence from international joint ventures. *Journal of Financial Economics* 73, 323–374.
- Doz, Y.L., Hamel, G., 1998. *Alliance Advantage: The Art of Creating Value through Partnering*. Harvard Business School Press, Boston, MA.
- Elfenbein, D.W., Lerner, J., 2003. Ownership and control rights in Internet portal alliances, 1995–1999. *RAND Journal of Economics* 34, 356–369.

- Esty, B.C., 1997. Organizational form and risk taking in the savings and loan industry. *Journal of Financial Economics* 44, 25–55.
- Frandsen, B.R., Frölich, M., Melly, B., 2012. Quantile treatment effects in the regression discontinuity design. *Journal of Econometrics* 168, 382–395.
- Fulghieri, P., Sevilir, M., 2009. Organization and financing of innovation, and the choice between corporate and independent venture capital. *Journal of Financial and Quantitative Analysis* 44, 1291–1321.
- Gomes-Casseres, B., Hagedoorn, J., Jaffe, A.B., 2006. Do alliances promote knowledge flows? *Journal of Financial Economics* 80, 5–33.
- Grossman, S.J., Hart, O.D., 1986. The costs and benefits of ownership: a theory of vertical and lateral integration. *Journal of Political Economy* 94, 691–719.
- Guedj, I., 2005. Ownership vs. contract: how vertical integration affects investment decisions in pharmaceutical R&D. *McCombs Research Paper Series FIN-01-06*.
- Guedj, I., Scharfstein, D., 2005. Organizational scope and investment: evidence from the drug development strategies and performance of biopharmaceutical firms. *NBER Working Paper No. 10933*.
- Gulati, R., 1995a. Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances. *Academy of Management Journal* 38, 85–112.
- Gulati, R., 1995b. Social structure and alliance formation patterns: a longitudinal analysis. *Administrative Science Quarterly* 40, 619–652.
- Habib, M.A., Mella-Barral, P., 2007. The role of knowhow acquisition in the formation and duration of joint ventures. *Review of Financial Studies* 20, 189–233.
- Habib, M.A., Mella-Barral, P., 2013. Skills, core capabilities, and the choice between merging, allying, and trading assets. *Journal of Mathematical Economics* 49, 31–48.
- Hackbarth, D., Mathews, R., Robinson, D., 2012. Capital structure, product market dynamics, and the boundaries of the firm. *Unpublished working paper*. University of Illinois at Urbana-Champaign, University of Maryland, and Duke University.
- Hart, O., Moore, J., 1990. Property rights and the nature of the firm. *Journal of Political Economy* 98, 1119–1158.
- Hendricks, K., Porter, R.H., 1992. Joint bidding in federal OCS auctions. *American Economic Review Papers and Proceedings* 82, 506–511.

- Hendricks, K., Porter, R.H., 1996. The timing and incidence of exploratory drilling on offshore wildcat tracts. *American Economic Review* 86, 388–407.
- Hendricks, K., Porter, R.H., Boudreau, B., 1987. Information, returns, and bidding behavior in OCS auctions: 1954–1969. *Journal of Industrial Economics* 35, 517–542.
- Hendricks, K., Porter, R.H., Tan, G., 2003. Bidding rings and the winner’s curse: the case of federal offshore oil and gas lease auctions. NBER Working Paper No. 9836.
- Hendricks, K., Porter, R.H., Tan, G., 2008. Bidding rings and the winner’s curse. *RAND Journal of Economics* 39, 1018–1041.
- Holmstrom, B., 1999. Managerial incentive problems: a dynamic perspective. *Review of Economic Studies* 66, 169–182.
- Imbens, G., Kalyanaraman, K., 2012. Optimal bandwidth choice for the regression discontinuity estimator. *Review of Economic Studies* 79, 933–959.
- Imbens, G.W., Lemieux, T., 2008. Regression discontinuity designs: a guide to practice. *Journal of Econometrics* 142, 615–635.
- Kent, D.H., 1991. Joint ventures and non-joint ventures: an empirical investigation. *Strategic Management Journal* 12, 387–393.
- Lafontaine, F., Slade, M., 2007. Vertical integration and firm boundaries: the evidence. *Journal of Economic Literature* 45, 629–685.
- Lerner, J., Malmendier, U., 2010. Contractibility and the design of research agreements. *American Economic Review* 100, 214–246.
- Lerner, J., Merges, R.P., 1998. The control of technology alliances: an empirical analysis of the biotechnology industry. *Journal of Industrial Economics* 46, 125–156.
- Lerner, J., Shane, H., Tsai, A., 2003. Do equity financing cycles matter? Evidence from biotechnology alliances. *Journal of Financial Economics* 67, 411–446.
- Masten, S.E., 1984. The organization of production: evidence from the aerospace industry. *Journal of Law and Economics* 27, 403–417.
- Mathews, R.D., 2006. Strategic alliances, equity stakes, and entry deterrence. *Journal of Financial Economics* 80, 35–79.
- McConnell, J.J., Nantell, T.J., 1985. Corporate combinations and common stock returns: the case of joint ventures. *Journal of Finance* 40, 519–536.

- Mitchell, W., Singh, K., 1996. Survival of businesses using collaborative relationships to commercialize complex goods. *Strategic Management Journal* 17, 169–195.
- Monteverde, K., Teece, D.J., 1982. Supplier switching costs and vertical integration in the automobile industry. *Bell Journal of Economics* 13, 206–213.
- Mowery, D.C., Oxley, J.E., Silverman, B.S., 1996. Strategic alliances and interfirm knowledge transfer. *Strategic Management Journal* 17(Winter Special Issue), 77–91.
- Palia, D., Ravid, S.A., Reisel, N., 2008. Choosing to cofinance: analysis of project-specific alliances in the movie industry. *Review of Financial Studies* 21, 483–511.
- Robinson, D.T., 2008. Strategic alliances and the boundaries of the firm. *Review of Financial Studies* 21, 649–681.
- Robinson, D.T., Stuart, T.E., 2007a. Network effects in the governance of strategic alliances. *Journal of Law, Economics, & Organization* 23, 242–273.
- Robinson, D.T., Stuart, T.E., 2007b. Financial contracting in biotech strategic alliances. *Journal of Law and Economics* 50, 559–595.
- Sampson, R.C., 2004. The cost of misaligned governance in R&D alliances. *Journal of Law, Economics, & Organization* 20, 484–526.
- Seru, A., forthcoming. Firm boundaries matter: evidence from conglomerates and R&D activity. *Journal of Financial Economics*.
- Sheen, A., 2009. Do public and private firms behave differently? An examination of investment in the chemical industry. Unpublished working paper. Harvard University.
- Stein, J.C., 1997. Internal capital markets and the competition for corporate resources. *Journal of Finance* 52, 111–133.
- Stuart, T.E., 2000. Interorganizational alliances and the performance of firms: a study of growth and innovation rates in a high-technology industry. *Strategic Management Journal* 21, 791–811.

Table 1
Firm bidding

This table summarizes the winning lease bids submitted by the ten most active firms in the data, as measured by the number of dollars contributed to winning bids. The sample is the 1,070 leases that were successfully sold to a solo firm or an alliance between 1954 and 1975, that received at least one solo bid and at least one alliance bid, and that have all data available. Bid values are adjusted to 1980 dollars.

	Number of solo bids	Number of alliance bids	Percent of bids as a solo firm	Dollars contributed to solo bids (millions)	Dollars contributed to alliance bids (millions)	Percent of dollars contributed as a solo firm
Exxon	61	27	69.3%	1,043.5	482.4	68.4%
Mobil	10	103	8.8%	157.5	1,310.2	10.7%
Texaco	36	62	36.7%	851.0	561.9	60.2%
Gulf	40	61	39.6%	506.5	742.3	40.6%
Shell	82	10	89.1%	723.3	154.9	82.4%
Tenneco	31	38	44.9%	427.7	388.3	52.4%
Getty	2	112	1.8%	34.4	723.6	4.5%
Sunoco	49	11	81.7%	695.7	16.9	97.6%
Chevron	35	53	39.8%	277.5	406.6	40.6%
Amoco	11	100	9.9%	55.4	617.6	8.2%
All firms	507	563	47.4%	5,735.0	11,777.0	32.7%

Table 2

Lease summary statistics

This table provides descriptive statistics at the lease level for the sample summarized in Table 1. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. *BHNPV* is the present discounted value of operating profits from a borehole, less the cost of drilling and equipping the borehole. *MEANBHNPV* is the mean of *BHNPV* over all boreholes drilled on a lease within five years of the lease start date. *LSNPV* is the net present value of all operating profits from boreholes drilled within 20 years of the lease start date, less the amount of the winning auction bid and the costs of drilling and equipping the boreholes. *MEANBHNPV* and *LSNPV* are winsorized at the 1st and 99th percentiles. Winning bids, *MEANBHNPV*, and *LSNPV* are adjusted to 1980 dollars.

	Mean	Std. dev.	1 st Pctile.	25 th Pctile.	Median	75 th Pctile.	99 th Pctile.	Non-missing obs.
Area (acres)	4,707	1,167	1,250	4,995	5,000	5,432	5,760	1,070
Winning bid per acre	3,558.87	5,229.90	63.53	606.94	1,688.08	4,055.84	24,915.03	1,070
<i>BIDRATIO</i>	0.465	0.273	0.002	0.221	0.468	0.680	0.991	1,070
Boreholes within 5 yrs.	5.45	8.92	0	1	2	6	45	1,070
Boreholes within 20 yrs.	9.73	15.66	0	1	4	11	74	1,070
<i>MEANBHNPV</i> (millions)	0.55	3.37	-3.54	-1.25	-0.71	1.11	14.88	924
<i>LSNPV</i> (millions)	6.19	55.99	-95.65	-13.58	-4.82	2.42	297.48	1,070

Table 3

Borehole summary statistics

This table provides descriptive statistics at the borehole level for boreholes drilled on a lease in the sample summarized in Table 1 within five years of the lease start date. *BHCOST* is the cost of drilling and equipping the borehole. *BHOP* is the present discounted value of operating profits from the borehole. *BHNPV* is *BHOP* minus *BHCOST*. *BHCOST*, *BHOP*, and *BHNPV* are winsorized at the 1st and 99th percentiles and are adjusted to 1980 dollars.

	Mean	Std. dev.	1 st Pctile.	25 th Pctile.	Median	75 th Pctile.	99 th Pctile.	Non-missing obs.
Feet drilled	9,285	3,347	693	7,226	9,402	11,669	16,393	5,833
<i>BHCOST</i> (millions)	1.46	0.89	0.12	0.84	1.29	1.85	4.66	5,833
<i>BHOP</i> (millions)	4.11	7.14	0	0	0.06	5.50	35.97	5,833
<i>BHNPV</i> (millions)	2.65	6.88	-3.44	-1.10	-0.49	3.88	33.90	5,833

Table 4

Mean borehole net present value for solo firms and alliances

This table presents the results of OLS regressions with mean borehole net present value $MEANBHNPV$, in millions of 1980 dollars, as the outcome variable. $BHNPV$ is the present discounted value of operating profits from the borehole, less the cost of drilling and equipping the borehole. $MEANBHNPV$ is the mean of $BHNPV$ over all boreholes drilled on a lease within five years of the lease start date. $MEANBHNPV$ is winsorized at the 1st and 99th percentiles. When a solo firm wins the lease auction, $BIDRATIO$ is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, $BIDRATIO$ is equal to the highest solo firm bid divided by the winning alliance bid. The sample is the sample summarized in Table 1, with the column headings indicating further sample restrictions based on $BIDRATIO$. Regressions using the optimal bandwidth are presented in the middle two columns. Bid values are adjusted to 1980 dollars. Standard errors robust to heteroskedasticity are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	$BIDRATIO > 0$		Optimal bandwidth: $BIDRATIO > 0.415$		$BIDRATIO > 0.708$	
Solo firm	-1.33*** (0.46)	-1.12** (0.45)	-1.64*** (0.62)	-1.61*** (0.62)	-2.53*** (0.85)	-2.18** (0.90)
log(Winning bid per acre)		0.47*** (0.09)		0.46*** (0.12)		0.22 (0.19)
Lease start year fixed effects	No	Yes	No	Yes	No	Yes
Separate linear functions of $BIDRATIO$, solo and alliance	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.011	0.086	0.023	0.122	0.044	0.151
Observations	924	924	505	505	206	206

Table 5

Lease net present value for solo firms and alliances

This table presents the results of OLS regressions with lease net present value *LSNPV*, in millions of 1980 dollars, as the outcome variable. *LSNPV* is the present discounted value of all operating profits from boreholes drilled within 20 years of the lease start date, less the amount of the winning auction bid and the costs of drilling and equipping the boreholes. *LSNPV* is winsorized at the 1st and 99th percentiles. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. The sample is the sample summarized in Table 1, with the column headings indicating further sample restrictions based on *BIDRATIO*. Regressions using the optimal bandwidth are presented in the middle two columns. Bid values are adjusted to 1980 dollars. Standard errors robust to heteroskedasticity are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	<i>BIDRATIO</i> > 0.186		Optimal bandwidth: <i>BIDRATIO</i> > 0.593		<i>BIDRATIO</i> > 0.797	
Solo firm	-17.96** (8.73)	-17.35* (8.88)	-30.76** (13.70)	-31.53** (14.90)	-46.58** (20.63)	-52.32** (23.42)
log(Winning bid per acre)		2.54 (1.94)		1.81 (2.57)		4.27 (4.73)
Lease start year fixed effects	No	Yes	No	Yes	No	Yes
Separate linear functions of <i>BIDRATIO</i> , solo and alliance	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.012	0.068	0.020	0.081	0.048	0.135
Observations	853	853	374	374	158	158

Table 6

Production, operating profits, and drilling costs for solo firms and alliances

This table separately examines the quantity of oil and gas production, the operating profits, and the drilling costs that are reflected in the two primary outcome measures. Borehole production is the present discounted value of the oil and gas produced by a borehole, applying oil and gas prices from 1960 to production in all years. Mean borehole production is the mean of borehole production over all boreholes drilled on a lease within five years of the lease start date. Mean borehole operating profits and mean borehole drilling costs represent the decomposition of mean borehole net present value *MEANBHNPV* into operating profit and drilling cost components. Lease production is the present discounted value of the oil and gas produced by boreholes drilled within 20 years of the lease start date, applying oil and gas prices from 1960 to production in all years. Lease operating profits and lease drilling costs represent the decomposition of lease net present value *LSNPV* into operating profit and drilling cost components, omitting the cost component due to the amount of the winning auction bid. Each of the six outcome variables is measured in millions of 1980 dollars and winsorized at the 1st and 99th percentiles. The sample is the sample summarized in Table 1. The OLS regressions each further limit the sample to observations within the optimal *BIDRATIO* bandwidth for the analogous regression with *MEANBHNPV* or *LSNPV* as the outcome variable. Bid values are adjusted to 1980 dollars. Standard errors robust to heteroskedasticity are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Outcome:	Mean borehole production	Mean borehole operating profits	Mean borehole drilling costs	Lease production	Lease operating profits	Lease drilling costs
Solo firm	-1.03*** (0.38)	-1.72*** (0.64)	-0.11 (0.12)	-27.91** (12.23)	-36.64* (19.41)	-6.66* (3.71)
log(Winning bid per acre)	0.43*** (0.09)	0.49*** (0.13)	0.02 (0.02)	12.59*** (2.31)	16.70*** (3.39)	5.10*** (0.73)
Lease start year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Separate linear functions of <i>BIDRATIO</i> , solo and alliance	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.152	0.165	0.402	0.182	0.161	0.217
Observations	505	505	505	374	374	374

Table 7

Distribution of experience among solo firms and alliances

This table reports the distribution of high-experience firms among solo firms and alliances that win lease auctions. For each set of auctions conducted simultaneously for leases in a given MMS area, a high-experience firm is defined as a firm for which the number of previously owned leases in the same MMS area is higher than the median among firms that participated in that set of auctions. A solo firm can be composed of zero or one high-experience firm, while an alliance can be composed of zero, one, or more high-experience firms. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. The first pair of columns reports the distribution of experience for all observations in the sample summarized in Table 1. The second and third pairs of columns report distributions of experience when the sample is further restricted to observations in the optimal *BIDRATIO* bandwidth regression samples when *MEANBHNPV* and *LSNPV*, respectively, are the outcome variables.

Number of high-experience firms	Primary analysis sample (<i>BIDRATIO</i> > 0)		Optimal bandwidth regression sample for <i>MEANBHNPV</i> (<i>BIDRATIO</i> > 0.415)		Optimal bandwidth regression sample for <i>LSNPV</i> (<i>BIDRATIO</i> > 0.593)	
	Solo firms	Alliances	Solo firms	Alliances	Solo firms	Alliances
0	307 (60.6%)	301 (53.5%)	144 (57.6%)	124 (48.6%)	116 (61.7%)	97 (52.2%)
1	200 (39.4%)	96 (17.1%)	106 (42.4%)	42 (16.5%)	72 (38.3%)	27 (14.5%)
2		86 (15.3%)		47 (18.4%)		37 (19.9%)
3		37 (6.6%)		18 (7.1%)		11 (5.9%)
4		41 (7.3%)		23 (9.0%)		14 (7.5%)
5		2 (0.4%)		1 (0.4%)		
Total	507	563	250	255	188	186

Table 8

Outcomes for solo firms and alliances accounting for experience

This table presents the results of OLS regressions where the outcome variable is mean borehole net present value *MEANBHNPV* (in millions of 1980 dollars) or lease net present value *LSNPV* (also in millions of 1980 dollars). Both outcome variables are winsorized at the 1st and 99th percentiles. The data for each regression are limited to solo firms and alliances composed of the indicated number of high-experience firms. For each set of auctions conducted simultaneously for leases in a given MMS area, a high-experience firm is defined as a firm for which the number of previously owned leases in the same MMS area is higher than the median among firms that participated in that set of auctions. All specifications control for the logarithm of the winning bid per acre (adjusted to 1980 dollars), lease start year fixed effects, and separate linear functions of *BIDRATIO* for solo firms and alliances. The *BIDRATIO* variable is recalculated for each regression sample. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. The sample is the sample summarized in Table 1. Each regression sample is further limited to observations within the *BIDRATIO* bandwidth that is optimal for the analogous regression that includes solo firms and alliances of all experience categories. Standard errors robust to heteroskedasticity are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Outcome:	<i>MEANBHNPV</i>				<i>LSNPV</i>			
Solo: high-exp. firm	0	1	0 or 1	0 or 1	0	1	0 or 1	0 or 1
Alliance: high-exp. firms	0	1	0 or 1	≥ 2	0	1	0 or 1	≥ 2
Solo firm	0.11 (0.68)	0.08 (2.77)	-0.68 (0.76)	-3.43*** (1.06)	-8.30 (11.20)	-6.05 (55.98)	-22.06 (17.29)	-46.15* (26.23)
R^2	0.176	0.298	0.124	0.215	0.132	0.578	0.110	0.127
Obs.	201	59	358	176	155	34	262	127

Table 9

Outcomes for solo firms and alliances controlling for firm effects

This table presents the results of OLS regressions where the outcome variable is mean borehole net present value *MEANBHNPV* (in millions of 1980 dollars) or lease net present value *LSNPV* (also in millions of 1980 dollars). Both outcome variables are winsorized at the 1st and 99th percentiles. Lead firm indicators are dummy variables that take a value of one when the relevant firm has the largest percentage ownership of the winning lease bid and a value of zero otherwise. Ties go to the firm with the highest number of dollars devoted to winning bids in the sample. Firm percent ownership variables are equal to the relevant firm's percentage ownership of the winning lease bid. Depending on the specification, either type of firm effect can be interacted with indicators for the lease starting in a given five-year window. All specifications control for the logarithm of the winning bid per acre (adjusted to 1980 dollars), lease start year fixed effects, and separate linear functions of *BIDRATIO* for solo firms and alliances. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. The sample is the sample summarized in Table 1. To aid in the estimation of the firm effects, the regression samples include all observations within twice the *BIDRATIO* bandwidth suggested by Imbens and Kalyanaraman (2012). Standard errors robust to heteroskedasticity are reported in parentheses, clustered at the lead firm level in specifications with lead firm indicators and clustered at the lead firm by five-year window level in specifications with lead firm indicators interacted with five-year window indicators. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Outcome:	<i>MEANBHNPV</i>				<i>LSNPV</i>			
Solo firm	-1.13** (0.43)	-1.03* (0.61)	-1.45*** (0.54)	-1.36** (0.66)	-22.28*** (7.56)	-14.93 (10.05)	-27.28*** (10.62)	-23.14* (12.85)
Lead firm indicators	Yes				Yes			
Lead firm × 5-year window	Yes				Yes			
Firm percent ownership	Yes				Yes			
Firm percent ownership × 5-year window					Yes			
R^2	0.146	0.189	0.227	0.288	0.143	0.199	0.191	0.308
Obs.	924	924	924	924	853	853	853	853

Table 10

Outcomes for solo firms and alliances controlling for financial variables

This table presents the results of OLS regressions where the outcome variable is mean borehole net present value *MEANBHNPV* (in millions of 1980 dollars) or lease net present value *LSNPV* (also in millions of 1980 dollars). Both outcome variables are winsorized at the 1st and 99th percentiles. Market leverage for a winning solo firm is the book value of the firm's debt divided by the sum of the book value of debt and the market value of the firm's equity, and market leverage for a winning alliance is the average market leverage of the constituent firms, weighted by percent ownership. Profitability for a winning solo firm is annual operating income after depreciation divided by total assets, and profitability for a winning alliance is the average profitability of the constituent firms, weighted by percent ownership. Available cash for a winning solo firm is the value of cash and short-term investments, and available cash for a winning alliance is the sum of available cash for the constituent firms. Total assets for a winning solo firm are taken directly from the firm's balance sheet, and total assets for a winning alliance are the sum of total assets for the constituent firms. Available cash and total assets are measured in millions of 1980 dollars. Market leverage, profitability, available cash, and total assets are winsorized at the 1st and 99th percentiles. All specifications control for the logarithm of the winning bid per acre (adjusted to 1980 dollars), lease start year fixed effects, and separate linear functions of *BIDRATIO* for solo firms and alliances. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. The sample is the sample summarized in Table 1, restricted to observations for which market leverage, profitability, available cash, and total assets are non-missing. The regression samples are further limited to observations within the optimal *BIDRATIO* bandwidth from the middle columns of Table 4 (when *MEANBHNPV* is the outcome variable) or the middle columns of Table 5 (when *LSNPV* is the outcome variable). Standard errors robust to heteroskedasticity are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Outcome: <i>MEANBHNPV</i>		<i>LSNPV</i>	
Solo firm	-2.11*** (0.82)	-1.96** (0.86)	-38.25* (19.62)	-35.83* (20.98)
Market leverage		0.35 (1.81)		53.31 (52.19)
Profitability		1.18 (6.43)		213.10 (151.10)
log(Available cash)		0.43 (0.57)		4.21 (10.54)
log(Total assets)		-0.36 (0.62)		-8.15 (11.54)
R^2	0.103	0.107	0.088	0.111
Obs.	339	339	248	248

Fig. 1. Distribution of leases by *BIDRATIO*, all leases. This figure displays the number of leases that fall in a given *BIDRATIO* bin for the sample summarized in Table 1. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid.

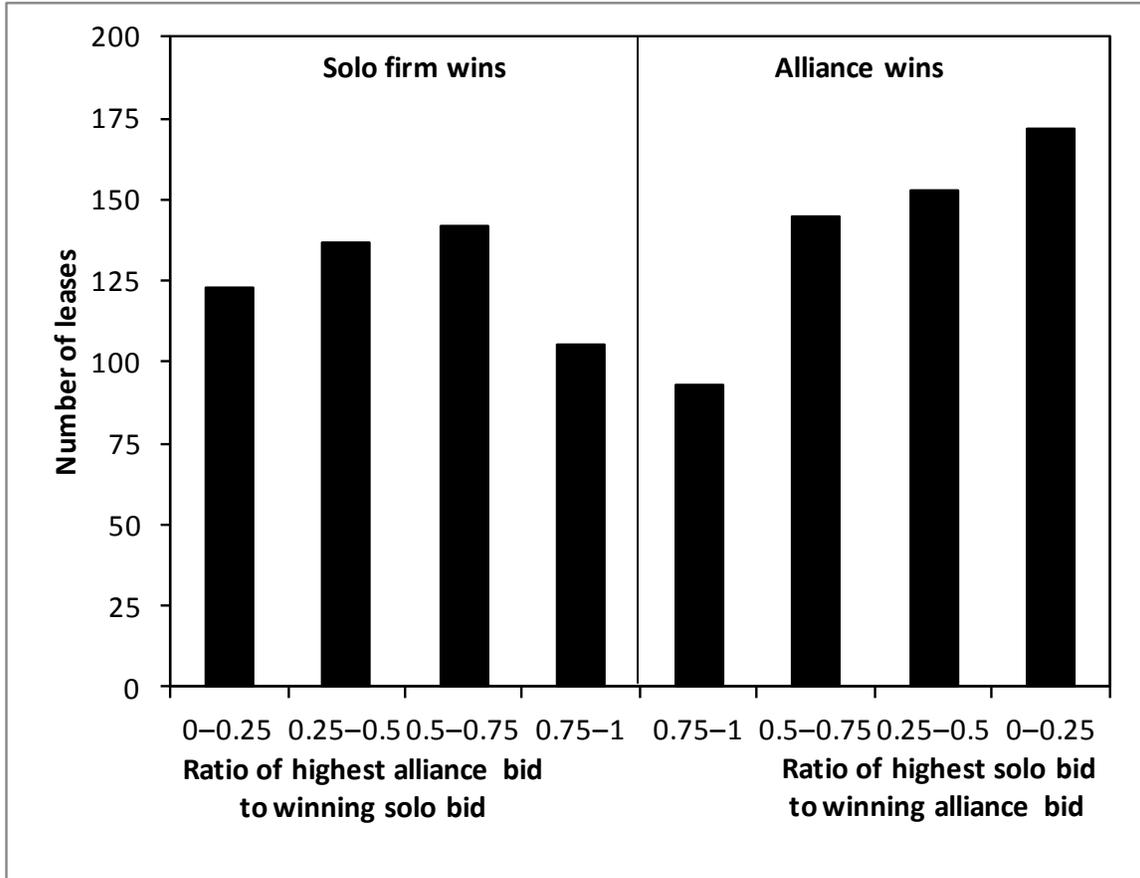


Fig. 2. Distribution of leases by *BIDRATIO*, $BIDRATIO > 0.5$. This figure displays the number of leases that fall in a given *BIDRATIO* bin for the sample summarized in Table 1, restricted to leases for which $BIDRATIO > 0.5$. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid.

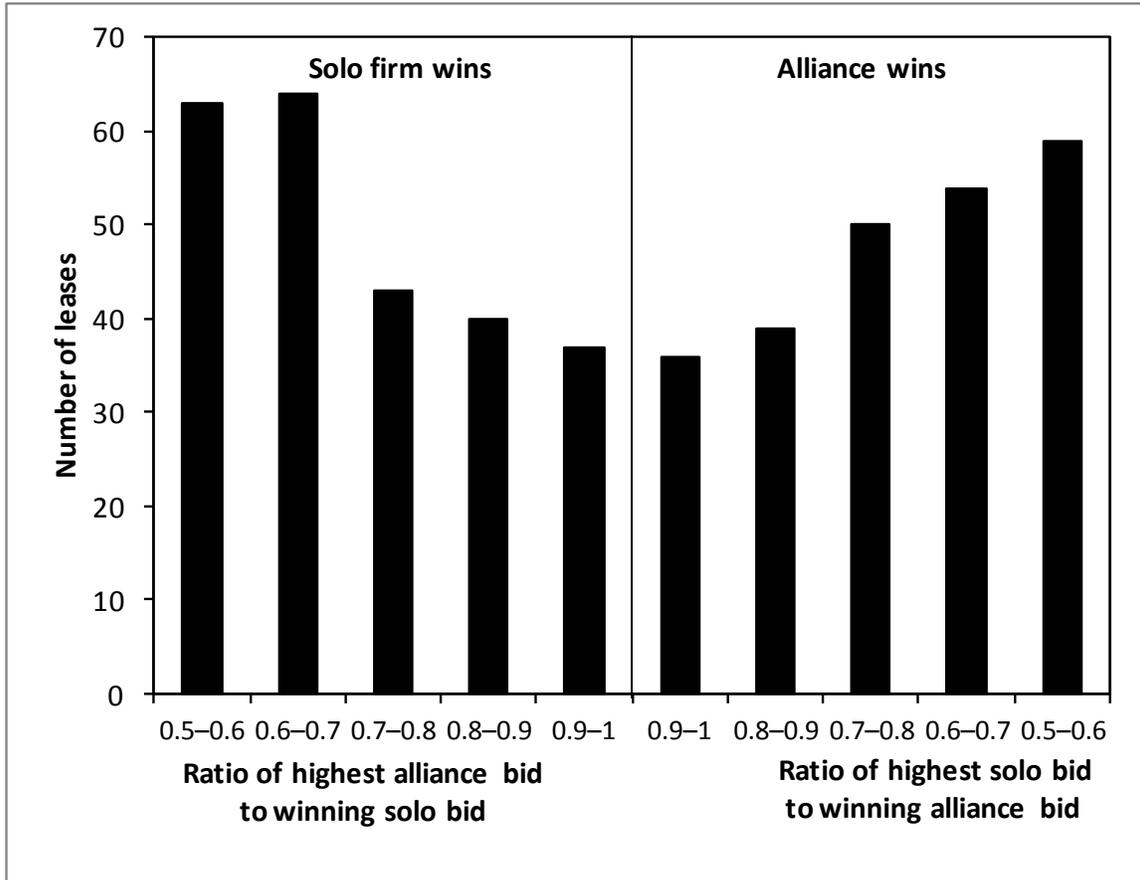


Fig. 3. Winning bids, all leases. This figure displays the mean winning bid for leases within a given *BIDRATIO* bin for the sample summarized in Table 1. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. The bars indicate two standard errors on either side of the mean.

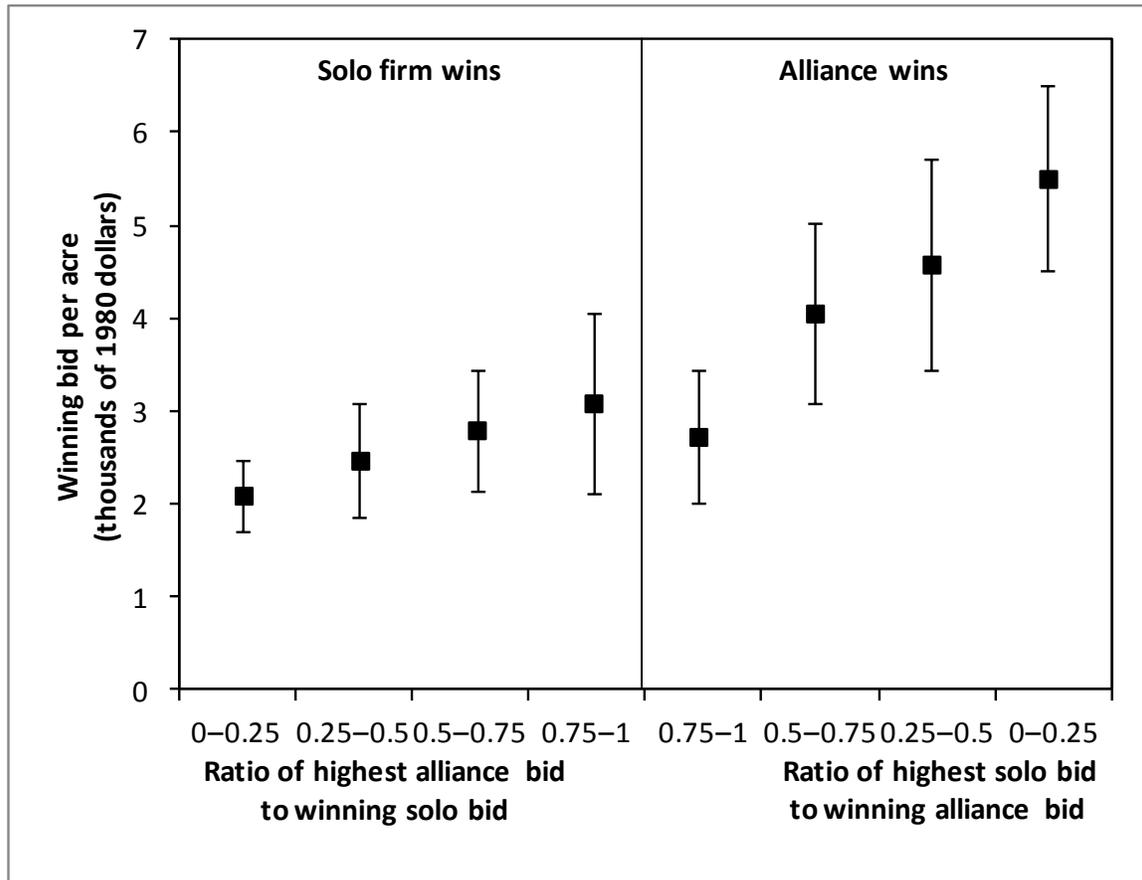


Fig. 4. Winning bids, $BIDRATIO > 0.5$. This figure displays the mean winning bid for leases within a given $BIDRATIO$ bin for the sample summarized in Table 1, restricted to leases for which $BIDRATIO > 0.5$. When a solo firm wins the lease auction, $BIDRATIO$ is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, $BIDRATIO$ is equal to the highest solo firm bid divided by the winning alliance bid. The bars indicate two standard errors on either side of the mean.

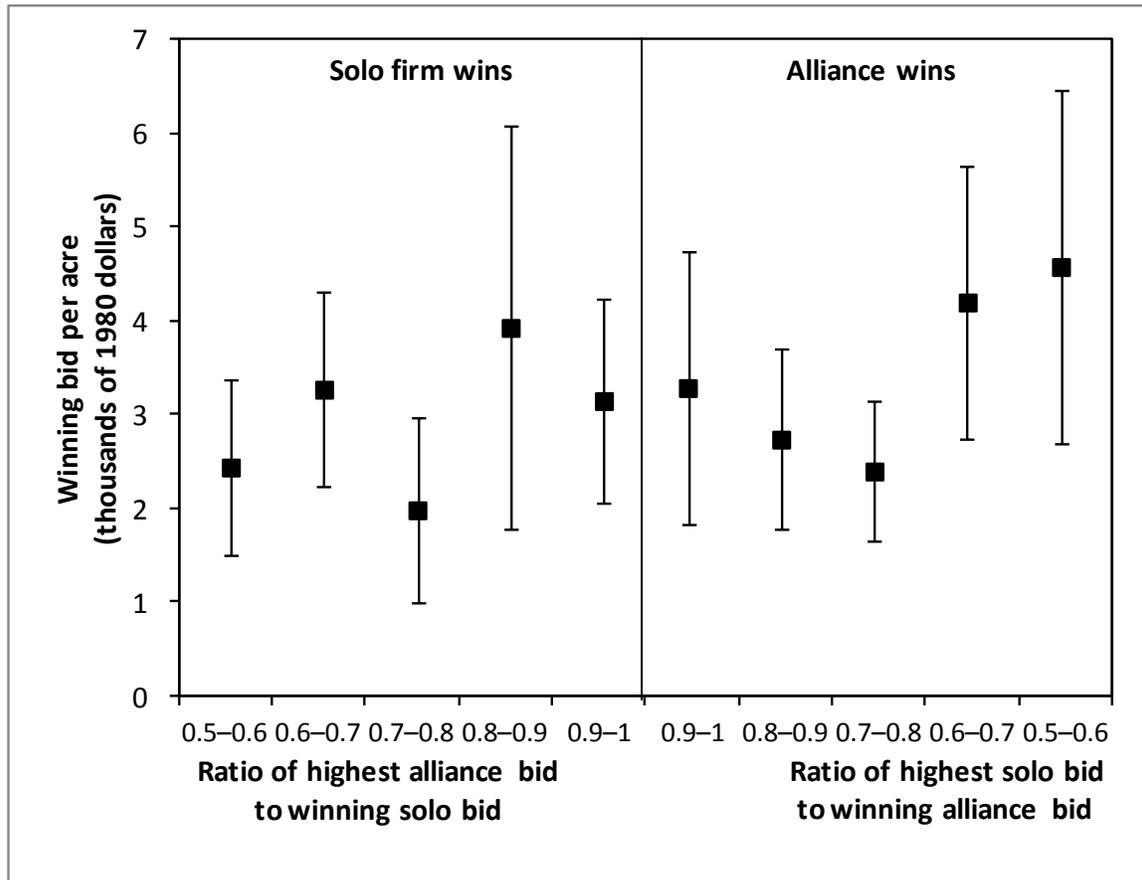


Fig. 5. Mean borehole net present value for solo firms and alliances. This figure displays the mean of $MEANBHNPV$ for leases within a given $BIDRATIO$ bin for the sample summarized in Table 1. $BHNPV$ is the present discounted value of operating profits from the borehole, less the cost of drilling and equipping the borehole. $MEANBHNPV$ is the mean of $BHNPV$ over all boreholes drilled on a lease within five years of the lease start date. $MEANBHNPV$ is winsorized at the 1st and 99th percentiles. When a solo firm wins the lease auction, $BIDRATIO$ is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, $BIDRATIO$ is equal to the highest solo firm bid divided by the winning alliance bid. The bars indicate two standard errors on either side of the mean. Predicted values are shown from regressions fitting separate linear functions of $BIDRATIO$ on each side of the discontinuity, using all leases in the sample or restricting the sample to leases with $BIDRATIO > 0.415$ (the optimal bandwidth).

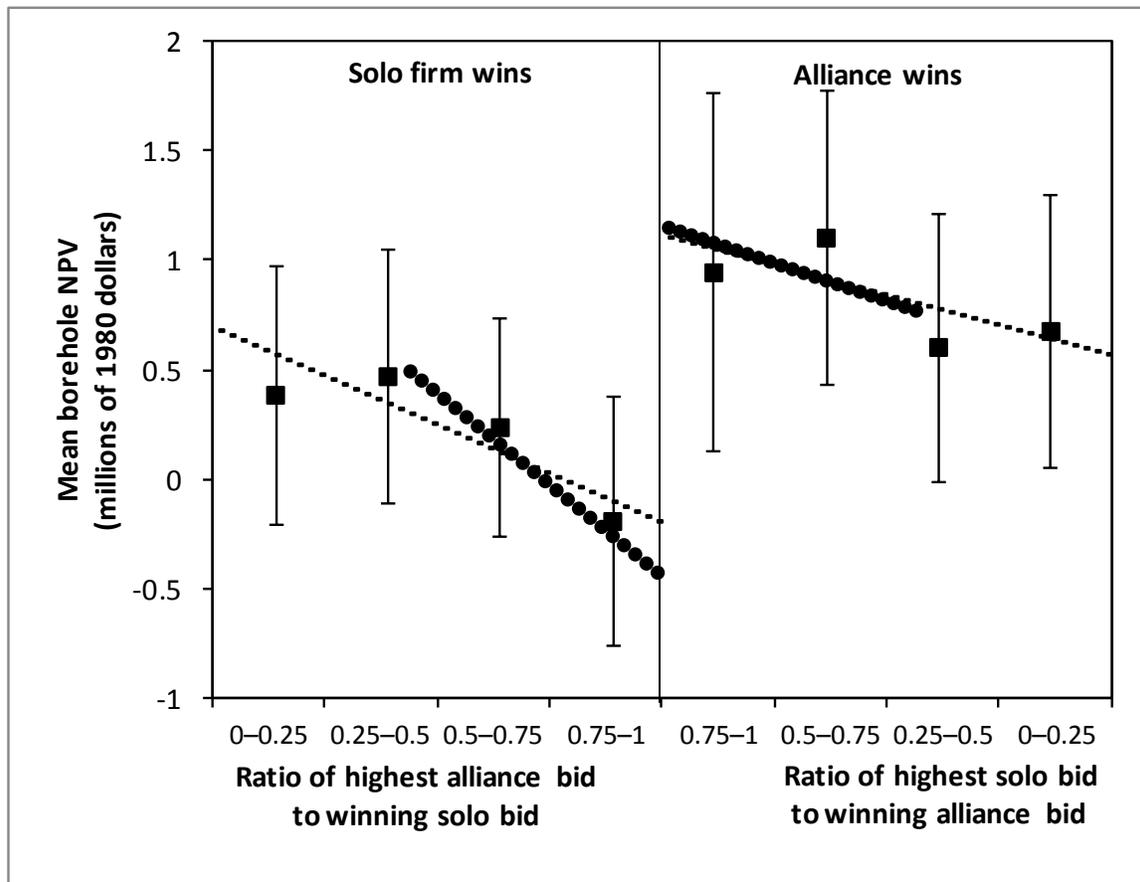


Fig. 6. Lease net present value for solo firms and alliances. This figure displays the mean of *LSNPV* for leases within a given *BIDRATIO* bin for the sample summarized in Table 1. *LSNPV* is the present discounted value of all operating profits from boreholes drilled within 20 years of the lease start date, less the amount of the winning auction bid and the costs of drilling and equipping the boreholes. *LSNPV* is winsorized at the 1st and 99th percentiles. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. The bars indicate two standard errors on either side of the mean. Predicted values are shown from regressions fitting separate linear functions of *BIDRATIO* on each side of the discontinuity, restricting the sample to leases with *BIDRATIO* > 0.186 or to leases with *BIDRATIO* > 0.593 (the optimal bandwidth).

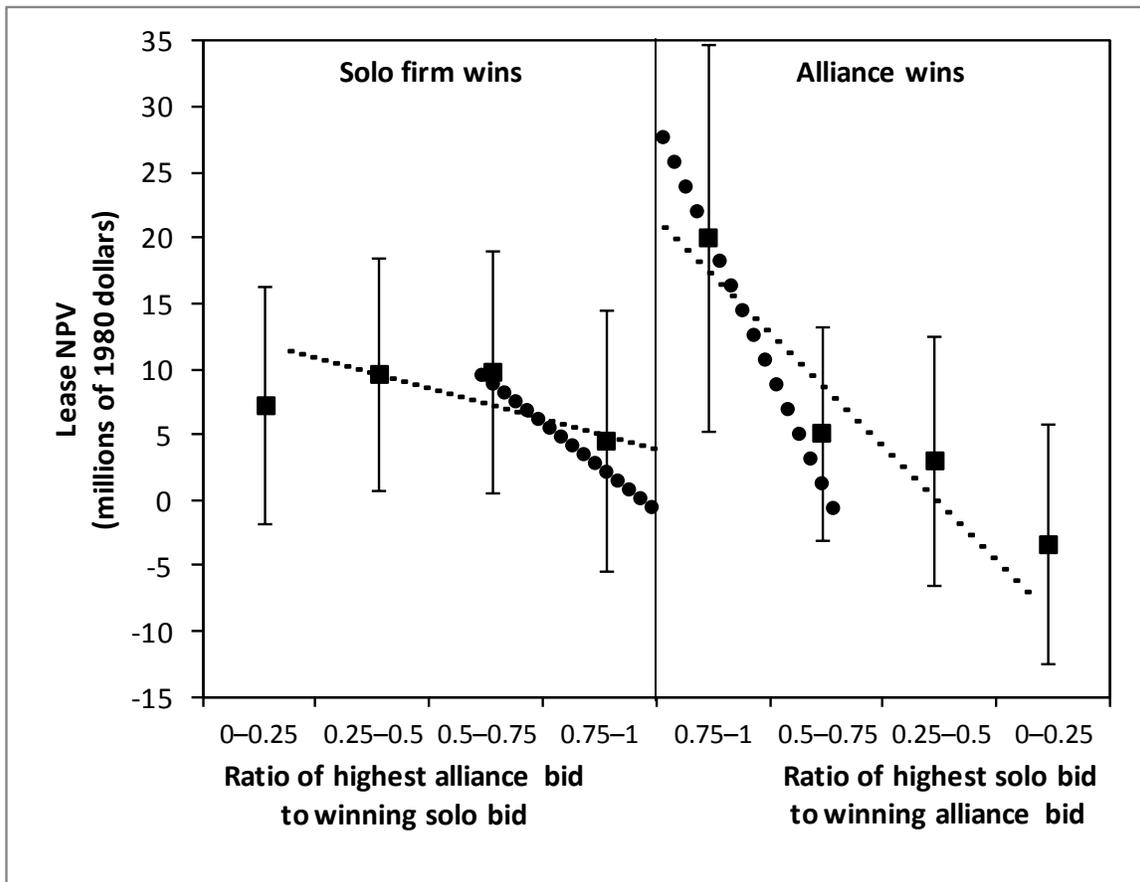


Fig. 7. Lease net present value for solo firms and alliances, quantile results. This figure displays the results from a series of regressions estimating the difference at the discontinuity between a given quantile of the *LSNPV* distribution for solo firms and the same quantile of the *LSNPV* distribution for alliances. The estimates are from quantile regressions fitting separate linear functions of *BIDRATIO* on each side of the discontinuity, restricting the sample to leases in the sample summarized in Table 1 for which *BIDRATIO* > 0.593 (the optimal bandwidth from the analogous estimation of the mean effect). *LSNPV* is the present discounted value of all operating profits from boreholes drilled within 20 years of the lease start date, less the amount of the winning auction bid and the costs of drilling and equipping the boreholes. *LSNPV* is winsorized at the 1st and 99th percentiles. When a solo firm wins the lease auction, *BIDRATIO* is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, *BIDRATIO* is equal to the highest solo firm bid divided by the winning alliance bid. The dotted lines indicate bias-corrected and accelerated bootstrapped 95% confidence intervals.

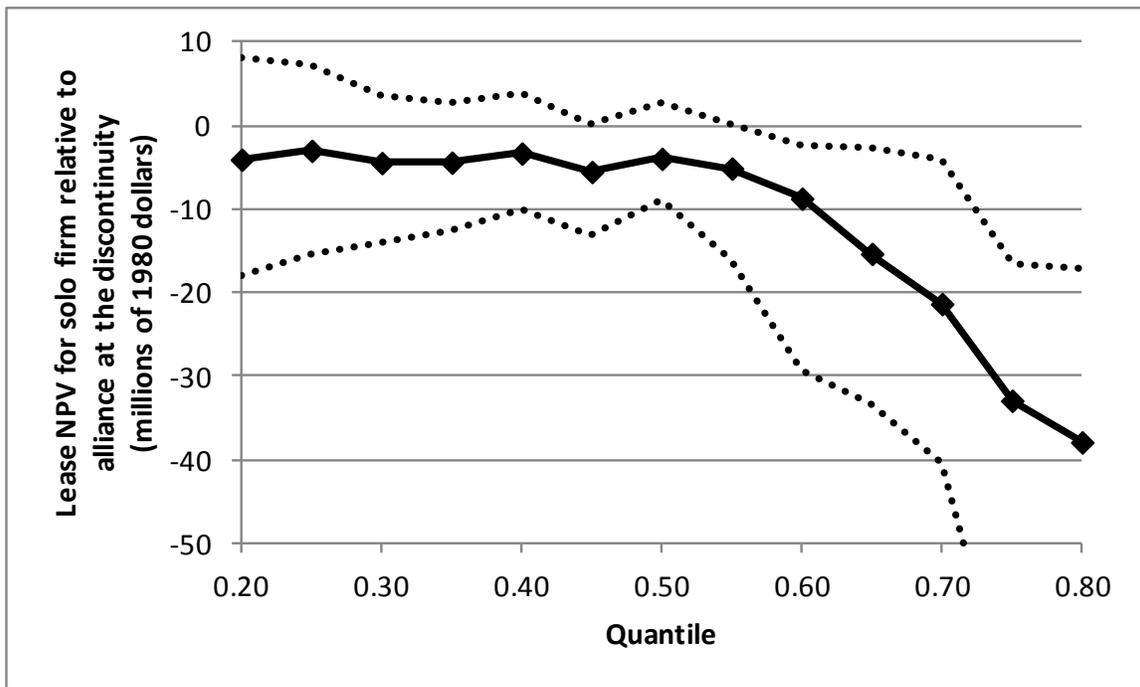


Fig. 8. Distributions of solo firm and alliance bids for leases close to the discontinuity and neighboring leases. This figure displays cumulative distribution functions for solo firm bids and for alliance bids on leases in the sample summarized in Table 1 for which $BIDRATIO > 0.9$ and neighboring leases. A neighboring lease is a lease located in a geographic block that contains a lease for which $BIDRATIO > 0.9$ or that touches a block containing a lease for which $BIDRATIO > 0.9$. When a solo firm wins the lease auction, $BIDRATIO$ is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, $BIDRATIO$ is equal to the highest solo firm bid divided by the winning alliance bid.

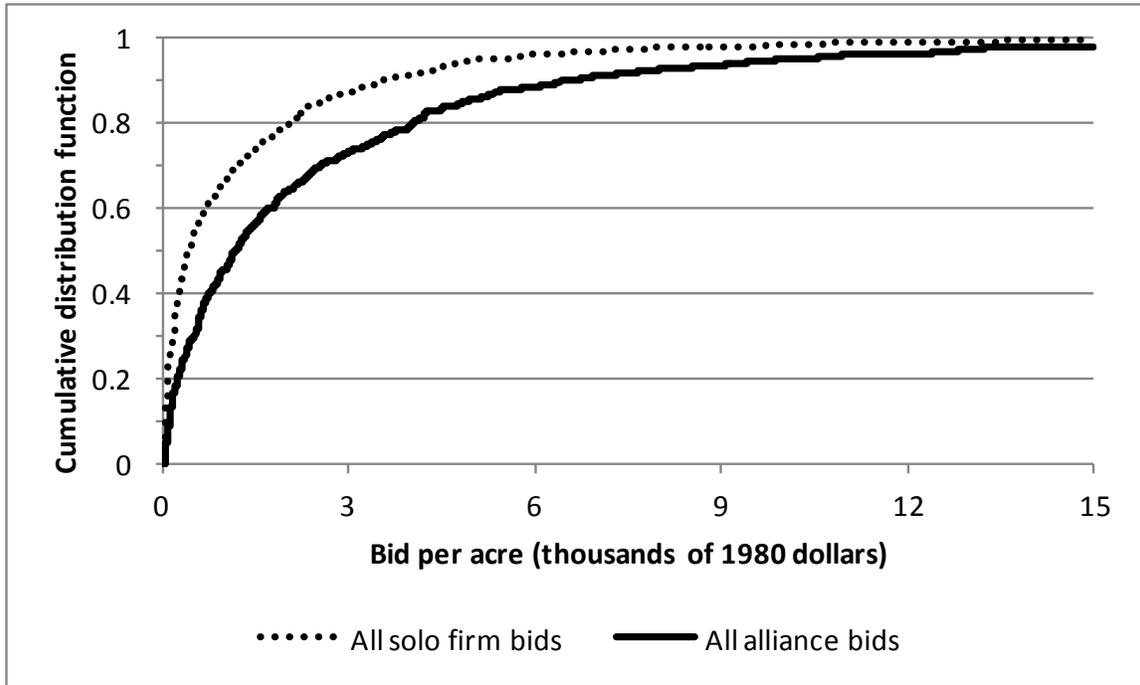


Fig. 9. Distributions of highest solo firm and alliance bids for leases close to the discontinuity and neighboring leases. This figure displays cumulative distribution functions for the highest solo firm bid and for the highest alliance bid on leases in the sample summarized in Table 1 for which $BIDRATIO > 0.9$ and neighboring leases. A neighboring lease is a lease located in a geographic block that contains a lease for which $BIDRATIO > 0.9$ or that touches a block containing a lease for which $BIDRATIO > 0.9$. When a solo firm wins the lease auction, $BIDRATIO$ is equal to the highest alliance bid divided by the winning solo firm bid; when an alliance wins the lease auction, $BIDRATIO$ is equal to the highest solo firm bid divided by the winning alliance bid. If no solo firms participate in a lease auction, the highest solo firm bid is coded as zero; if no alliances participate in a lease auction, the highest alliance bid is coded as zero.

