Title

Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch

Author

KELLOGG, RYAN M

Publication Date

2007-11-13

CSEM WP 173

Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch

Ryan Kellogg

November 2007

This paper is part of the Center for the Study of Energy Markets (CSEM) Working Paper Series. CSEM is a program of the University of California Energy Institute, a multi-campus research unit of the University of California located on the Berkeley campus.

2547 Channing Way
Berkeley, California 94720-5180
www.ucei.org

Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch

Ryan Kellogg*

November 6, 2007

Abstract

Production in many industries, such as construction and heavy manufacturing, relies on inputs from both lead firms and contractors. These firms' joint productivity often hinges on their ability to share information and coordinate activities, suggesting that they have strong incentives to learn about each other's personnel, procedures, and expertise. This learning differs from standard learning-by-doing in that it is relationship-specific: its benefits are not appropriable outside the relationship in which the learning takes place. In this paper, I empirically examine the importance of relationship-specific learning using high-frequency data from oil and gas drilling. I find that the joint productivity of a lead firm and its drilling contractor is enhanced significantly as they accumulate experience working together. This result is robust to other relationship specificities. I also find that firms appear to recognize the benefits of joint experience: controlling for other specificities, lead firms are more likely to work with contractors with which they have substantial prior experience than those with which they have worked relatively little.

---

*Department of Agricultural and Resource Economics at UC Berkeley, and University of California Energy Institute (UCEI), kellogg@are.berkeley.edu. I am grateful for financial support from the OpenLink Fund at the Coleman Fung Risk Management Center. This paper has benefited from discussions with Maximilian Auffhammer, Severin Borenstein, Meghan Busse, Michael Greenstone, Jeffrey Perloff, and numerous seminar participants. Tay Feder provided excellent research assistance.

1. Introduction

The outsourcing of productive activity is common in many industries, ranging from construction and heavy manufacturing to white-collar business services. Construction projects, for example, regularly involve multiple contractors and sub-contractors working under a lead, general contractor. Productive efficiency requires not only that each firm be adept at its own set of tasks, but that the firms effectively coordinate their activities and share information. Firms responsible for project design must align their efforts with the firms that actually carry out the construction, and general contractors must efficiently plan the overlapping activities of subcontractors to avoid delays. Consider Boeing's recent launch of the 787 Dreamliner passenger jet, which involved nearly 30 firms contracted directly with Boeing, as well as countless additional subcontractors and suppliers. Collaboration amongst these firms has been central to the jet's development and production—one manager publicly commented that "interpersonal communication skills and building relationships have become more important than ever" (*Managing Automation* 2007).

The productivity benefits of inter-firm collaboration suggest that lead firms and contractors have strong incentives to learn about each other's personnel, procedures, and expertise. Such learning can take place as the firms accumulate experience working together. For example, an accounting firm may improve the speed with which it prepares a client's quarterly reports as its employees become familiar with the client's personnel and accounts.[1] This learning differs from standard learning-by-doing in that it is *relationship-specific*. The accounting firm will generally not be able to apply its knowledge of one client to augment its productivity with a different client.

This paper empirically investigates relationship-specific learning through an examination of the oil and gas drilling industry. I find strong evidence that the accumulation of experience

[1] Groysberg (2001) and Huckman and Pisano (2006) document evidence that an individual's job performance is influenced by knowledge and skills that are specific to the firm in which he or she works. In this example, however, employees are acquiring knowledge and skills specific to their firm's clients, rather than to their own firm.

specific to pairs of firms working together plays an important role in improving drilling performance. Such experience effects have not been identified in prior studies of learning-by-doing, despite the fact that many of the industries often explored in learning studies involve substantial contracting and are likely subject to relationship-specific learning. Wright (1936), Alchian (1963), and Benkard (2000) investigate aircraft manufacturing; Joskow and Rose (1985) investigate power plant construction; and Argote *et al*. (1990) and Thornton and Thompson (2001) examine World War II shipbuilding. All of these authors confirm that performance in these industries improves with overall production experience, but do not assess whether some of the observed productivity increases are relationship-specific.

Relationship-specific learning is important because its economic implications extend beyond its direct impact on firms' productivity. In particular, it provides firms with an incentive to maintain stable contracting relationships. Relationship-specific learning creates intellectual capital that is not appropriable across different contracting partners; thus, a lead firm is likely to prefer a contractor with which it has substantial experience over one with which it has worked relatively little. An objective of this paper is to examine empirically the persistence of firm relationships in the drilling industry, and test whether observed contracting patterns are consistent with firms' recognition of the benefits of joint experience.

Though not explicitly examined in this paper, relationship-specific learning also has implications that relate to the literature on transactions costs, contractual completeness, and the boundaries of the firm.[2] An important strand of this research, pioneered by Williamson (1975, 1979, 1985) and Klein *et al*. (1978), emphasizes the role of relationship-specific investments in driving long-term contracting and vertical integration. These authors argue that the rents generated by such investments can lead to opportunistic bargaining problems when the *ex ante* contract does not specify how the rents are to be divided. Moreover, Williamson (1979) argues

[2] Ronald Coase (1937) was the first to emphasize the roles of control and transactions costs in defining firm boundaries. See Whinston (2003) and Gibbons (2005) for surveys of the literature on the theory of the firm, and Lafontaine and Slade (2007) for a survey of related empirical work.

that the relevant relationship-specific investments may be physical or intellectual in nature. Thus, relationship-specific learning and the intellectual capital it generates may play a role in promoting long-term contracts and vertical integration.

The oil and gas drilling industry is well-suited for a study of relationship-specific learning for several reasons. First, there exists vertical separation between two types of firms that are involved in drilling: production companies ("producers") and drilling companies.[3] Producers—for example, ExxonMobil and Chevron—are responsible for the technical design and planning of wells to be drilled, but do not actually drill wells themselves. Drilling is instead outsourced to drilling companies that own and staff drilling rigs. Second, learning is an important source of productivity growth in this industry. Drilling cost-efficiency, driven almost entirely by minimizing the time required to drill each well, requires the technical optimization of drilling procedures as well as effective teamwork between producer personnel and the rig crew.

I have obtained well-level industry contracting and performance data, covering nearly 20,000 wells drilled over 1991-2005, with which I track drilling efficiency for producers, rigs, and producer-rig pairs. My primary finding is that not only do producers and rigs learn from their own experience, they also benefit from relationship-specific learning. Specifically, a rig that works with only one producer will, on average, benefit from productivity improvements twice as large as those of a rig that frequently changes producers. I verify that this result is not driven by other specificities between producers and rigs that might cause certain firm pairings to drill more effectively than others.

I also find that the pattern by which producer-rig pairs are formed and broken is consistent with firms' recognition of relationship-specific learning's benefits. In particular, when a producer has contracted with multiple rigs in the same county, and releases one of its rigs to another producer, the rig most likely to be released is that with the least producer-specific

[3] As discussed further in section 2, this vertical separation allows producers to drill wells with greater spatial and temporal flexibility.

experience. Evidence suggests that this pattern is driven by learning effects rather than other producer-rig specificities.

Beyond these primary results, I also test for the presence of cross-producer learning spillovers in the drilling industry. Such spillovers are generally important, as macroeconomic theory indicates that they are important drivers of economic growth (Parente 1994, Jovanovic and Nyarko 1996). However, when I analyze the productivity of producers working side-by-side in the same oilfield, I find little evidence of spillover effects. This result stands in contrast to other studies that identify modest learning spillovers in semiconductor manufacturing and shipbuilding (Irwin and Klenow 1994, Thornton and Thompson 2001).

The remainder of the paper is organized as follows: section 2 provides general background information on the oil and gas drilling industry, and section 3 discusses industry mechanisms and incentives for learning-by-doing. Section 4 describes the data used in this study. Section 5 presents the empirical framework and estimation results for learning-by-doing by production companies, omitting the influence of the rigs they hire. This provides a baseline for section 6, which presents evidence of relationship-specific learning. Section 7 discusses relationship persistence between producers and rigs, and section 8 offers concluding comments.

2. The Onshore Oil and Gas Drilling Industry

2.1 Production companies and the drilling problem

Oil and gas reserves are found in distinct geologic formations known as fields that lie beneath the earth's surface, and the mission of a production company is to extract these reserves for processing and sale. To operate in any given field, a producer must first obtain leases from the holders of that field's mineral rights. A lease typically grants a right to operate in only a small part of a field, and most fields are operated and drilled by several distinct producers holding different leases.[4] In contrast to leases in the federal offshore continental shelf, which have been studied by Hendricks and Porter (1988) among others, there is no centralized process in Texas by which producers obtain leases. Instead, producers proactively approach the holders of mineral rights on the land that they wish to explore, who may then negotiate lease terms or organize a competitive bidding process.

A field's reserves are typically buried under many layers of rock that do not contain oil or gas. The objective of drilling a well is to penetrate these overlying rock layers to reach the oil and gas in the field. While the geology within any given field is quite homogenous, there is significant geologic variation across fields, particularly with regards to the depth at which they are buried. Some fields are found as shallow as 1,000 feet and can be drilled in a few days, while others are more than 20,000 feet deep and can require several months of drilling. The geologic composition of the rock that must be drilled through also varies considerably. Multiple layers of sandstone, shale, and limestone may be encountered as a well is drilled from the surface to its targeted depth in the field, and the types of rock encountered in one area will generally not be the same as those encountered elsewhere.

Wells fall into two broad categories. "Wildcats" are those wells that are drilled into a previously unexplored field, and their goal is to assess whether the field will actually be productive. "Development" wells, on the other hand, are drilled into fields in which previously drilled wells already exist, and their goal is to enhance field production. Most wells are vertical holes; however, horizontal and directional wells are sometimes drilled when surface features make a vertical well impossible, or when doing so will improve the well's ultimate production of oil and gas. These wells are technically more complex than vertical wells and may require substantially more time to drill.

[4] Leaseholding producers within a field may sometimes "unitize" their holdings by pooling them together, agreeing on ownership shares in the pooled unit, and naming one of the producers as the unit operator. See Wiggins and Libecap (1985) for a discussion of the economics of unitization.

Even though producers do not physically drill their own wells, they do design wells and write drilling procedures. This arrangement is a response to the fact that the optimal drilling program for any well is a function of the specific geologic features of the field in which it is drilled. Producers typically have more geologic information than do drillers, due to their knowledge from seismic imaging and previously drilled wells, and are therefore better placed to make these engineering decisions.[5]

2.2 Rigs and contracting

The actual drilling of wells is carried out by drilling companies, which own drilling rigs and employ drilling crews. A typical onshore drilling rig is pictured in figure 1. Its primary features are a tall derrick, which allows pipe to be drawn in and out of the well, and a motor that spins the drill pipe and drill bit during drilling. The size of this equipment determines a rig's "depth rating," the depth to which the rig is recommended to drill. Apart from this depth rating, rigs generally do not have field or producer-specific characteristics. The exceptions to this rule are recently-built or refurbished rigs carrying computer equipment that eases the drilling of horizontal and directional wells. Because this equipment specificity may confound my analysis of relationship-specific learning, I ultimately omit horizontal and directional wells from the data.

Rigs are mobile and can easily change locations within a field; however, moves of more than 50 miles typically require several days and result in the charging of fees to the producer requesting the move. When under contract, rigs operate 24 hours per day and 7 days per week, rotating crews in three 8-hour shifts. Industry participants have indicated that, while the average employment tenure of a rig crewman is approximately one year, the rig foreman usually stays with a rig for much longer, and tenures longer than five years are not uncommon.

It is natural to ask why this industry is vertically separated, particularly given the relationship-specific learning effects identified in this paper. The answer lies in the spatial and

[5] Very small producers, which drill infrequently and may not have engineering resources, sometimes outsource the planning and design function to the driller, particularly if the driller has experience in the same field.

temporal variation with which producers drill wells. The drilling activity of any producer fluctuates with oil and gas prices, and with its success in finding new fields. Successful wildcats and development wells often lead to additional drilling, while dry holes do not. The mobility and non-specificity of rigs allow them to smooth these fluctuations in drilling requirements across nearby producers. This smoothing minimizes overall rig capacity requirements, as well as rig transportation and mobilization costs, without requiring the producers to contract directly with each other.

Producers typically contract with rigs for the drilling of one well at a time, though a producer and rig will write a multi-well contract when the producer is confident in its future demand for wells. Producers initiate the contracting process by issuing a request for quotation (RFQ) from drilling companies with rigs in the vicinity of the proposed well. The RFQ contains detailed technical specifications regarding the well to be drilled, including the well's total depth, the diameters and lengths of steel well casing strings to be installed in the well, and the density of the "drilling mud" to be pumped through the borehole during drilling. The driller then includes in its bid, along with price, the identities of the rig and crew it proposes to drill the well.

The RFQ also specifies which of two standard contract types will be used: "dayrate" or "footage." In a dayrate contract, the drilling company provides a rig and crew to drill the well to the producer's specifications, charging it a daily payment for the rig's services. The producer is represented on the rig by one of its personnel, known as the "company man," who directs the rig's operations, typically in consultation with the rig's foreman. In a footage contract, the rig is compensated at a rate set in dollars per foot drilled. This contract type is equivalent to a fixed-price contract since the well's depth is specified in advance in the RFQ. The producer may or may not place a company man on the rig. If present, he may monitor the rig's activities and consult with the rig foreman on drilling decisions, but has no direct contractual authority.[6]

[6] While the determinants of contract type have been examined by Corts (2004) in the offshore drilling industry, the effect of contract type on performance remains an unanswered empirical question. Research into this issue using the data discussed here requires controls for the endogenous choice of contract type and is therefore left as a topic for future research. However, in section 6.3, I do address the possibility that the learning analysis presented in this paper is affected by contract choice.

3. Productivity and Learning-by-Doing in Onshore Drilling

In this paper, I use the time necessary to drill a well as the measure of drilling productivity. While this approach is necessitated by the fact that I lack well-level cost data, it parallels the way producers and engineers actually view drilling efficiency, and is arguably superior to using cost data were such information available. In practice, drilling engineers achieve cost savings almost entirely by reducing the time necessary to drill wells. Given dayrates that typically exceed $10,000 per day, saving a day's worth of rig time is well worth the efforts of producers' engineering teams.[7] In addition, given a particular well and rig, there is little scope for substitution between drilling time and labor or capital. Rigs always work 24 hours per day and 7 days per week, and adding crew members cannot make the drilling bit turn more quickly. Most capital drilling inputs, such as the casing and tubing that are installed in the well, are fixed functions of the well's depth. For these reasons, learning curve case studies in the petroleum engineering literature use drilling time as their performance metric, even though the authors typically have access to detailed cost data. Brett and Millheim (1986) argue that the drilling time metric is actually superior to a cost metric, since cost data are polluted by inconsistent accounting methods and variations in materials prices and rig rates. Moreover, rig rates are likely to be endogenous: the prices charged by rigs rise during periods of high drilling activity, which will create spurious correlation between drilling cost and experience.

Producers have ample scope to learn from prior drilling experiences. Each new penetration into a field yields information regarding both the field's geology and which drilling procedures work well in that geology. Many drilling decisions such as choice of mud density and bit selection depend critically on the types of rock encountered. Thus, learning on the part of field producers is technical in nature and tends to be very field-specific. This learning is well-recognized within the drilling industry, and several published engineering case studies have

[7] In the case of footage contracts, efficient well design, backed by historically low drilling times, can be used to obtain lower bids from drillers. Moreover, cutting days from a drilling program reduces a producer's use of secondary contractors, such as those supplying fuel and water to the rig, yielding additional cost savings.

documented how lessons learned from experience have been applied to reduce drilling times. See, for example, Brett and Millheim (1986) and Adeleye *et al*. (2004).

Because rigs are not involved in well design and planning, rig-level learning is less technical in nature than is producer-level learning. Instead, rigs' learning effects are based on developments in each crew member's skills and on improved teamwork. These performance improvements generate clear economic benefits to a rig and its drilling company when the rig contracts on a footage basis, but indirect benefits exist under dayrate contracts as well. Industry participants have indicated that rig reputations are well-known by producers, and that rigs that are known to have effective, experienced crews can command a dayrate premium over other rigs. Also, because the producer's company man is present on the rig on a dayrate contract, he can observe the efforts of the rig foreman and crew very easily. In an environment in which repeat contracting is very common, this observability of effort can create implicit incentives to perform well, as shown theoretically by Corts (2007).

Finally, and of principal importance in this paper, there exists scope for relationship-specific learning between a producer and a rig. Many specific mechanisms of such learning are possible. The rig's crew may become familiar with the producer's particular drilling procedures, or the company man may improve his knowledge of the capabilities of the rig and its crew. In addition, the ability to rapidly solve drilling problems—for example, a loss in the circulation of drilling mud or the sticking of pipe in the wellbore—is an important determinant of drilling efficiency. Industry participants have indicated that these problems are more easily solved if the company man and rig foreman have developed a working relationship that allows them to collaborate effectively.

4. Data

The central empirical challenge of this paper is to separate the impact of relationship-specific learning from the effects of other, non-relationship-specific forms of learning. My approach uses two datasets of drilling activity in Texas. I obtained the first of these from the Texas Railroad Commission (TRRC), Texas's regulator of oil and gas drilling activity. These data consist of well-level records of every well drilled in the state from 1977-2005. Each observation identifies the field and county in which the well was drilled, and the identity of the producer that drilled the well. I take the number of days required to drill each well as the difference between the well's completion date and the date drilling began. This latter date was not regularly recorded until 1991: only 42% of observations have a drilling time prior to this date, compared to 90% afterwards. I therefore focus my analysis on the 1991-2005 period.[8]

Each record in the TRRC data also indicates the well's depth, whether the well was drilled to produce oil or gas, and whether the well is vertical, horizontal, or directional. Because horizontal and directional wells are typically best drilled with specialized rigs, I omit these wells, comprising 21% of the data, from my analysis.

The TRRC data do not include the identities of the drilling rigs that drilled each well. I therefore obtained information on rig activity from Smith Bits (SB). Smith Bits is a manufacturer of drilling bits, and its field sales force issues weekly reports on all onshore rig activity in North America. These reports give each rig's location, by county, on every Friday from 1989 to 2005, and also provide the identity of the production company to which the rig is contracted. Each observation also includes the depth of the well being drilled, the rig's depth rating, and whether the well is being drilled for oil or for gas.[9] Unlike the TRRC data, the unit of observation in the

[8] While the TRRC asks producers to report date drilling began for all their wells, this reporting is not rigorously enforced. Beyond these missing data, 3.7% of the observations from 1991-2005 have drilling times that are clearly erroneous or technically infeasible. I drop wells with drilling times that are negative, wells with drilling times greater than 180 days, and wells that are more than 3,000 feet deep and (implausibly) reported to have been drilled in a single day. The incidence of these observations and those with missing drilling times does not appear to be correlated with the experience variables that I ultimately use in my analysis.

[9] Unfortunately, I do not observe whether the rig is on a one-well or multi-well contract, or the price charged.

SB data is a rig-week, and I do not observe individual wells. Thus, if the SB data indicate that a particular rig spends three consecutive weeks working for the same producer in the same county, I cannot discern, without additional information, whether that rig has drilled three very quick wells or one long well.

My empirical analysis of relationship-specific learning requires a well-level dataset in which each observation reports the well's drilling time, location, producer, and drilling rig. I construct this dataset by merging the SB rig location data into the TRRC's well-level drilling records. Unfortunately, a large fraction of wells in the TRRC data cannot be matched to rig information in the SB data. Match failures occur for three reasons. First, some wells in shallow fields are drilled in less than one week and may therefore not be drilled on a Friday. Such wells, comprising 6% of the TRRC dataset, have no corresponding record in the SB data and are therefore impossible to match.[10]

Second, 51% of the remaining TRRC wells do not match because the producer names in the TRRC data do not always agree with the producer names in the SB data. Often, two names are similar only in part, and it is difficult to discern whether the two names do in fact point to the same firm. I use information on firm addresses, officer names, and drilling frequency to carefully match some similar names; however, I leave many ambiguous cases unmatched to avoid the risk of matching firms that are, in fact, distinct.

Finally, some non-unique matches occur when a producer employs multiple drilling rigs simultaneously in the same county. Because the SB data do not contain field or well information, I am unable to distinguish which rig is drilling which well in such cases. While I am able to use information on well depth and well type to match some of these wells to their rigs, there are other cases in which there is no way to confidently match the data. Rather than guess, I drop all wells that cannot be matched uniquely. This reduces the dataset by a further 39%.

[10] In section 6.3, I verify that the selective removal of wells drilled in less than one week does not impact the analysis.

Following the match, I drop all fields, producers, and rigs for which there is only one observation, since tracking learning for such entities is not possible. This procedure leaves a matched dataset of 19,174 wells, amounting to approximately one-quarter of the original TRRC sample. These wells are spread over 1,606 fields, 779 producers, and 1,334 rigs. As indicated in table 1, there is a large variance in drilling activity across these entities. For example, in some fields I observe only two wells, while in others I observe hundreds. Table 1 also indicates variance in the number of producers working within any field: some fields are drilled by only one producer and others are drilled by more than ten.

Figure 2 illustrates the relation between drilling time and depth in the sample. Very shallow wells that are a few thousand feet deep may be drilled in less than a week, whereas wells deeper than 15,000 feet can require several months of drilling. The sample average drilling time is 23.7 days, the average well depth is 9,040 feet, and 90% of the data lie between 3,200 feet and 14,000 feet. Summary statistics for depth, drilling time, and well type are presented in table 2.

5. Empirical Analysis: Learning by Field Producers

I begin the empirical analysis by examining the effect of producers' experience on their drilling productivity, omitting the influence of their relationships with rigs. This analysis follows existing learning-by-doing studies that investigate lead firm productivity but do not incorporate contractor relationships into their analysis. In section 6, I examine how the results presented here are affected by taking relationship-specific learning into account.

5.1 Empirical framework

The empirical framework is designed to capture how a producer's drilling times in a given field are influenced by that producer's field-specific experience, controlling for heterogeneities across fields and producers, and for technological change over time. The reference case specification is given by (1) below.

$$log(DrillTime_{fpt}) = f(Experience_{fpt}) + \gamma_f + \delta_p + \nu_t + \boldsymbol{\varphi X_{fpt}} + \varepsilon_{fpt} \qquad (1)$$

The dependent variable for each well—identified by its field *f*, producer *p*, and date of completion *t*—is the logarithm of the well's drilling time. The explanatory variable of primary interest is producer *p*'s experience in field *f* at time *t*, denoted by $Experience_{fpt}$. Setting aside for the moment the precise definition of $Experience_{fpt}$, the remaining variables included in (1) are vital in controlling for other factors that influence drilling time. These variables include:

γ_f: Field fixed effects

δ_p: Producer fixed effects

ν_t: Year fixed effects

$\boldsymbol{X_{fpt}}$: Well depth; flags for oil vs. gas well; flag for dry hole; month-of-year fixed effects

The field fixed effects γ_f control for the substantial heterogeneity in drilling conditions across fields. The producer fixed effects δ_p control for heterogeneity in drilling skill, and the year fixed effects ν_t control for industry-wide technological change.[11] I include variables $\boldsymbol{X_{fpt}}$ for well type (oil vs. gas and productive vs. dry) and well depth to control for remaining within-field heterogeneity. Month-of-year fixed effects control for seasonal variations in drilling time that may arise from changes in weather. The disturbance ε_{fpt} represents the presence or lack of drilling problems on each well, and is presumed to be heteroskedastic and correlated across wells drilled within the same field.

Given these fixed effects and controls, the effect of experience on drilling time is identified through variations in each producer's drilling activity within a field. There exist numerous sources of such variation. Increases in oil and gas prices will increase the number of wells drilled, though not uniformly across fields and producers (some fields will be on the margin at a given price while others will not be). Drilling may also be spurred by lease

[11] In alternative specifications, I use a polynomial function of time to control for technological change. Doing so does not substantially affect the estimated results.

acquisitions, discovery of new fields, or the identification of unexploited reserves in existing fields (through seismic imaging technology, for example).

I define $Experience_{fpt}$ as the number of wells drilled by producer p in field f during the two years prior to date t, including the well completed at t.[12] I calculate this variable using the original TRRC dataset rather than the smaller dataset generated by the match of the TRRC data to the SB data. Were I to instead use this smaller dataset, I would vastly understate each producer's experience, as the matched data include only one-quarter of the original TRRC observations. In addition, variations in the retention of data across fields and producers would add noise to the calculation, which could cause attenuation bias in the estimation of (1).

I measure experience using the number of wells drilled within the past two years rather than the total cumulative number of wells drilled because the majority of the fields in the dataset were discovered prior to the start of the sample.[13] I therefore have no means to calculate cumulative experience for wells in these fields. Even so, it is not clear that experience gained many years before time t is relevant to a producer's expertise at t. Studies by Argote *et al*. (1990) and Benkard (2000) have demonstrated that experience effects decay with time as learning is "forgotten," supporting the importance of recent experience in determining productivity. In this paper's appendix, I discuss evidence of forgetting effects in the drilling industry.[14]

Measuring experience using drilling activity over a fixed time period does come with a cost: it is likely to create simultaneity bias that will cause an estimate of (1) to exaggerate the learning effect. Sometimes, producers will hire a rig to drill a series of wells one right after another. In such cases, the number of wells drilled within any fixed time period will be inversely related to the number of days required to drill each well. For example, a producer that can drill a

[12] The inclusion of the well completed at t implies that all wells in the dataset have at least one unit of experience and avoids taking a logarithm of zero in a log-log specification of learning.

[13] The choice of two years is a compromise between capturing the tenures of rig crews and rig foremen. I discuss the results' robustness to measurements of experience using periods shorter or longer than two years in section 6.3.

[14] In the appendix, I also indicate that some of the forgetting effects that I observe may reflect losses of intellectual capital associated with changes in producers' drilling rigs. That is, I find that a mechanism of forgetting in the drilling industry may be the learning specificities that are the primary focus of this paper.

well in 20 days will drill 36 wells over two years, whereas a producer that requires only 15 days to drill a well will drill 49 wells over two years. Thus, decreases in drilling time due to learning may actually cause an increase in the number of wells drilled. This simultaneity will cause a spurious negative correlation between drilling time and $Experience_{fpt}$ in (1), exaggerating the estimated learning effect.

I address this problem by instrumenting for $Experience_{fpt}$ using an alternative measure of recent experience: the total number of days during which producer *p* actively drilled in field *f* during the two years prior to *t*. This variable is not subject to the simultaneity problem: when wells are drilled back-to-back, the total number of drilling days will remain roughly constant as the drilling time per well decreases and the number of wells drilled increases. Moreover, this instrument is clearly correlated, both intuitively and empirically, with $Experience_{fpt}$, the number of wells drilled.[15]

To capture within-firm learning spillovers across fields, an alternative specification of (1) includes a variable that counts the number of wells recently drilled by producer *p* in fields other than field *f*. I also estimate cross-firm spillovers using a variable that counts the number of wells recently drilled in field *f* by producers other than producer *p*. This variable is similar to those used by Irwin and Klenow (1994) and Thornton and Thompson (2001) in their studies of cross-firm learning spillovers. Summary statistics for all experience variables are presented in table 3.

5.2 Estimation results

To begin, I estimate (1) without allowing for within-firm or cross-firm learning spillovers. Most studies of learning-by-doing model learning curves with a log-log functional form, which in this setting implies that $f(Experience_{fpt}) = \beta \cdot \log(Experience_{fpt})$. Before taking this approach, I estimate (1) flexibly by fitting a cubic spline to $f(Experience_{fpt})$. The results are

[15] As an alternative to the instrumental variable strategy, I could use days of drilling directly as the measure of experience in (1). However, learning by producers is technical and driven by the geologic information gained with each penetration rather than the accumulation of days of experience. Thus, measuring experience using drilling days will lead to measurement error and attenuation bias. Indeed, using this measure of experience directly in (1) leads to estimated learning effects that are approximately half as large as those reported in the reference case.

plotted in figure 3. Drilling times are estimated to decrease by about 15% over the first 50 wells drilled by a field producer, and then stay relatively constant over the remaining wells. This productivity improvement is modest relative to gains found in other studies of learning-by-doing, reflecting the fact that drilling technology for onshore vertical wells is quite mature. Learning in this setting does not come from applying new technology or designing new products but from learning which existing drilling techniques are best applied to each geologic problem.

Also plotted on figure 3 is the result of estimating the log-log functional form. The point estimate of β is -0.036 with a clustered standard error of 0.005.[16] This specification closely matches the spline for wells with fewer than 100 units of experience, as highlighted in figure 4. It does over-predict productivity gains at very high levels of experience, though it remains within the 95% confidence interval of the spline estimate. Wells with greater than 100 units of experience carry little weight in the estimation as they constitute less than 3% of the data. Dropping these wells does not significantly affect the log-log estimate: β is estimated to increase in magnitude only to -0.037.

Table 4 displays the full set of estimated coefficients for the log-log specification plotted in figure 3, which I now refer to as the reference case. The estimated coefficients on the control variables generally agree with intuition. In particular, deeper wells require more drilling time than shallow wells. Dry holes require more drilling time than do productive wells, perhaps reflecting time spent trying to coax the well to flow. The year fixed effects indicate the presence of some technological change over the sample period, as drilling times decrease by approximately 15% from 1991 to 2005.

The second column of table 5 reports the results of estimating (1) without instrumenting for experience. These results agree with the anticipated direction of bias: the uninstrumented

[16] All standard errors presented in this section and in section 6 use a robust variance estimator that is clustered at the field level (Arellano 1987, Wooldridge 2003). This estimator allows for both heteroskedasticity and within-field correlation in the disturbance ε_{fpt}. Clustering on producer or on rig yields nearly identical results.

learning rate is larger than that of the reference case. A Hausman test strongly rejects the exogeneity of experience in this specification, with a p-value less than 0.001.

Could the reference case estimate be driven by producers' selection of fields in which to drill? If a producer is particularly good at drilling wells in certain types of fields—for reasons other than learning—and drills in such fields more frequently than in others, then a spurious negative correlation between experience and drilling time will be manifest in the data. I control for this possibility by adding fixed effects for field-producer interactions to specification (1). Doing so actually increases the magnitude of the estimated learning rate, as shown in column III of table 5. This result indicates that the observed experience effects are driven by learning rather than the matching of producers to fields for which they have specific drilling expertise.

In regression IV, I examine the importance of learning spillovers. Producers' experience in other fields appears to improve their productivity, though the magnitude of this effect is approximately one-third that associated with producers' field-specific learning. On the other hand, the estimate of cross-firm learning spillovers is small and statistically insignificant. This result contrasts with those of Irwin and Klenow (1994) and Thornton and Thompson (2001), which identify modest spillovers in the semiconductor and shipbuilding industries, respectively. Drilling industry participants have indicated that the lack of spillovers may be due to the competitive nature of common pool resource extraction. When multiple producers operate in the same field, an increase in production by one firm may deplete the resource in a way that adversely affects the production of the other firms. Thus, producers may be unwilling to aid each other by sharing their drilling procedures.

6. Empirical Analysis: Rigs and Relationship-Specific Learning

The analysis presented thus far has been that of a traditional learning-by-doing study in that it has omitted the contributions of contractors—drilling rigs—to the learning process. The learning estimates presented in tables 4 and 5 attribute all learning effects to producers without acknowledging the possibility that a share of this learning could be driven by rig or relationship-specific experience. This section takes advantage of producer-rig contracting data to examine rig and relationship-specific learning, and to assess the degree to which the previous section's results misattributed learning effects solely to producers.

6.1 Empirical framework

I augment specification (1) with variables that track rig and relationship-specific experience, and with rig fixed effects that control for rig heterogeneity. The new reference case specification is given by (2) below.

$$log(DrillTime_{fprt}) = f(\boldsymbol{Experience_{fprt}}) + \gamma_f + \delta_p + v_t + \eta_r + \boldsymbol{\varphi X_{fpt}} + \varepsilon_{fprt} \tag{2}$$

In (2), each well's field, producer, rig, and date are indexed by the subscripts f, p, r, and t, respectively. Rig fixed effects are denoted by η_r, while γ_f, δ_p, and v_t denote field, producer, and year fixed effects as in (1). $\boldsymbol{Experience_{fprt}}$ is now a vector of experience variables, and I expand $f(\boldsymbol{Experience_{fprt}})$ per (3) below.

$$f(\boldsymbol{Experience_{fprt}}) = \beta_1 \cdot log(Experience_{fpt}) + \beta_2 \cdot log(Experience_{-fpt}) \tag{3}$$
$$+ \beta_3 \cdot log(Experience_{f\text{-}pt}) + \beta_4 \cdot log(Experience_{-prt}) + \beta_5 \cdot log(Experience_{prt})$$

The first three terms of the expansion denote the three dimensions of experience examined in section 5: the experience of producer p in field f, the experience of producer p in other fields, and the experience of other producers in field f. As was the case in section 5, experience by producer p in field f is instrumented using the total number of drilling days accumulated by producer p in field f within the past two years.

The fourth term in (3) represents the experience of rig r with producers other than p, and the fifth term captures relationship-specific learning by measuring the experience of rig r with producer p. I calculate these two variables using the SB dataset before it is matched to the TRRC data as this avoids understating each rig's experience. Because the SB data track rig-weeks rather than wells, I define a rig's experience with producer p at time t to be the number of weeks it was actively drilling for p within the two years prior to t. Summary statistics for $Experience_{-prt}$ and $Experience_{prt}$ are indicated in the upper section of table 6.[17]

Relationship-specific learning is important to the extent that β_5 is more negative than β_4. There are two reasons to be concerned that the estimates of these coefficients may be confounded. First, a rig's experience with a particular producer is likely to be positively correlated with its experience in a particular field. Thus, my estimates of β_4 and β_5 may be confounded if rigs learn from field-specific experience. I control for this possibility by decomposing each rig's experience into the following field-specific and non-field-specific components:[18]

(1) $Experience_{-f-prt}$: experience with producers other than p in fields other than f

(2) $Experience_{f-prt}$: experience with producers other than p in field f

(3) $Experience_{-fprt}$: experience with producer p in fields other than f

(4) $Experience_{fprt}$: experience with producer p in field f

[17] The well being drilled by rig r at time t is included in $Experience_{-prt}$ and $Experience_{prt}$, as well as the four field-specific and non-field-specific experience variables. This avoids taking logarithms of zero.

[18] This decomposition of experience is complicated by the fact that the SB data do not contain field identifiers. Thus, even though I can identify each rig's field location for each matched observation, I cannot do so for every week in which each rig is active. I solve this problem in two steps. First, within the matched data, I find the fraction of wells drilled by each rig within the past two years that were in the same field as the rig's current field. Then, I multiply this fraction by the number of weeks the rig has been active during the past two years, taken from the SB data. This calculation yields an estimate of the number of weeks of experience that each rig has accumulated in its current field within the past two years. Other methods of calculating each rig's field-specific experience yield estimates of learning similar to those presented in this paper. For example, I have calculated experience based on the assumption that whenever I observe a rig in a particular field, it stays in that field until I observe it elsewhere.

Summary statistics for these dimensions of experience are indicated in the lower section of table 6. I use these variables to test whether a rig's experience specific to both its current field and current producer has a greater effect on drilling time than does its experience specific only to its current field.

The estimate of relationship-specific learning may also be confounded if some rigs have intrinsic producer-specific capabilities that are not generated from learning. If such rigs work more frequently with producers with which they are most compatible, there will be a spurious negative correlation between rig-producer joint experience and drilling times. Though industry participants have indicated that rigs are generally non-specific pieces of equipment (with the exception of their depth ratings), I verify that rig-producer specificities do not drive the estimate of (2) by adding fixed effects for rig-producer pairs to the specification. These fixed effects control for the possibility that some rigs may be more effective when working for some producers than for others, independently of learning effects.[19] Identification of relationship-specific learning therefore only comes from variations in joint experience within each rig-producer pair.

6.2 Estimation results

Column I of table 7 reports the results of estimating relationship-specific learning per equation (2). The estimated coefficient on log$(Experience_{prt})$—joint experience between a rig and a producer—is -0.021, statistically significant at the 1% level. This result implies that a rig that works with the same producer over one year can expect to decrease its drilling times by 8%. However, the estimated coefficient on log$(Experience_{-prt})$ is only -0.010, implying that a rig that frequently changes producers during a year can expect to decrease its drilling times by only 4%.

[19] Even when rig-producer fixed effects are included in the specification, I am, strictly speaking, only estimating a relationship-specific learning rate for those rig-producer pairs that I actually observe in the data. If producers are more likely to work with rigs with which they anticipate having steep learning curves, then hypothetical learning rates for the rig-producer pairs that I don't observe could be lower than the learning rate I estimate here. Short of being able to run a randomized experiment, there exist no plausible means to estimate an "average" learning rate over all possible rig-producer pairs. However, it is not clear that such a learning rate is actually a parameter of greater economic interest than the learning rate for relationships that actually occur in the industry.

Thus, on average, rigs with stable contracting relationships improve their productivity twice as quickly as rigs that frequently change contracting partners. Moreover, the difference between the coefficients on log*(Experience$_{prt}$)* and log*(Experience$_{-prt}$)* driving this result is statistically significant: an F-test rejects pooling with a p-value of 0.036.

In addition, the point estimate corresponding to learning by field producers is only -0.022 in this specification. This point estimate is substantially lower in magnitude than was reported in column IV of table 5, when the impact of producers' relationships with rigs was not considered in the regression. Thus, investigating learning using experience variables specific only to producers overestimates the contribution of their stand-alone experience to observed productivity improvements.

Column II of table 7 examines whether relationship-specific learning is driven by individual rigs or by rigs' parent entities: drilling companies. I estimate that the effect of a producer's joint experience with a drilling company, conditioned on that producer's experience with a particular rig, is nearly zero and statistically insignificant. This result indicates that relationship-specific learning is driven by the local interactions between a producer's personnel and a rig's crew rather than the producer's interactions with drilling company management.

This specification also indicates that a rig's performance is likely to be adversely affected if the other rigs owned by its drilling company increase their activity. The estimated effect of experience by the same drilling company but a different rig is positive and statistically significant at the 10% level. This result is puzzling at first, but can be explained by changes in the quality of rig crews. Industry participants have indicated that increases in activity by a drilling company are often accompanied by the hiring of "green" crew members that are particularly poorly skilled. Thus, when a drilling company activates many rigs and its overall firm experience rapidly increases, the performance of its rigs may suffer.

When I add fixed effects for rig-producer pairs to the specification, the estimated effect of rig-producer joint experience is not significantly affected, as shown in column III. This result suggests that rigs either do not have producer-specific capabilities (apart from those created through learning), or that they do not tend to work more frequently with producers with which they are most compatible. Were this tendency true, the estimate of relationship-specific learning would be lower in magnitude in column III than in column I.[20]

Column IV of table 7 decomposes each rig's experience into field-specific and non-field-specific components. These results lend additional support to the importance of relationship-specific learning: a rig's experience specific to both its current field and producer has a much stronger effect on drilling time than does experience specific only to its field. The point estimates imply that a rig that works for the same producer in the same field for a year can expect a 14% increase in drilling productivity. However, were the rig to then switch producers, its productivity would on average be only 3% larger than that of a rig with no experience at all. The difference in the estimated coefficients on log$(Experience_{f\text{-}prt})$ and log$(Experience_{fprt})$ that drives this result is statistically significant at the 1% level. In addition, regression V indicates that the significant effect of log$(Experience_{fprt})$ is robust to fixed effects for rig-producer-field interactions.

Column IV also highlights that rigs' acquisition of field-specific knowledge is important: a rig's experience that is specific to its current producer but not its current field is estimated to have little effect on productivity.[21] This result is likely to be driven by variations in producers' drilling procedures or personnel across fields.

6.3 Robustness of results

Exclusion of shallow wells from the data: As part of the merge process, some wells that were drilled in less than one week were dropped from the sample because they could not be matched to records in the Smith Bits data. Although these wells constitute only 6% of the overall

[20] With rig-producer fixed effects, the estimated effect of log($Experience_{\text{-}prt}$) becomes nearly zero and statistically insignificant. This change in the estimate occurs because the variation in log($Experience_{\text{-}prt}$) is largely accounted for by the fixed effects: each change in a rig's producer generates a new fixed effect. Thus, this result should not be interpreted as evidence that $Experience_{\text{-}prt}$ does not significantly affect productivity.

[21] A test for pooling of the three coefficients on log($Experience_{\text{-}f\text{-}prt}$), log($Experience_{f\text{-}prt}$), and log($Experience_{\text{-}fprt}$) fails to reject pooling with a p-value of 0.51. However, an F-test does reject that these three coefficients are jointly equal to zero with a p-value of 0.07.

population, it is possible that this selection on the dependent variable may bias the results. I address this concern by estimating (2) with data only for wells that are at least 8000 feet deep (12,654 observations). Such wells are essentially impossible to drill in less than one week, even under ideal conditions. Estimation with this sub-sample therefore neutralizes the potential selection problem.

Estimation results, presented in column I of table 8, are very similar to those obtained from the full sample (table 7, column I). Performance improves twice as quickly for rigs that have stable contracting relationships as for those that do not. The difference between the coefficients on log$(Experience_{-prt})$ and log$(Experience_{prt})$ is statistically significant with a p-value of 0.068. Column II indicates that, when experience is decomposed into field-specific and non-field-specific components, a rig's experience specific to both its current field and producer remains significantly more important than experience significant only to its field.

Flexible functional form: The reference case restricts the effect of $Experience_{fpt}$ to act through a log-log function, which may distort the estimated effects of rigs' producer-specific and field-specific experience. In columns III and IV of table 8, I allow $Experience_{fpt}$ to enter the regression as a spline, to verify that the finding of relationship-specific learning is not driven by functional form. The estimated joint experience effects are not substantially affected by this change. This result is consistent with the tight fit of the log-log model to the spline model, shown in figures 3 and 4.

Contract type: Could variation in the type of contract used by producers and rigs be influencing the reported results? The analysis to this point has not taken firms' choice of dayrate or footage contract into account. These contract types occur with roughly equal frequency in the data: 47.8% of contracts are dayrate while the remainder are footage. While both the producer and the rig will typically have indirect performance incentives under either contract type, the choice between dayrate and footage affects which firm has the direct incentive. Contract choice could therefore influence both the level of drilling productivity and the rate of productivity improvement. In particular, the decreases in drilling times I observe for producer-rig pairs could be driven by changes in contract type as producers and rigs accumulate experience together.

While it is tempting to control for contract type directly by including contract dummies in equation (2), this approach is plagued by the endogeneity of contract choice. Corts (2004) investigated the determinants of contract type in the offshore drilling industry, and his findings are consistent with what I find in my sample of onshore data. For example, dayrate contracts are relatively more common for deep wells and horizontal wells that require sophisticated well designs and typically involve substantial geological risks that are out of the rig's control.

Corts (2004) also finds that dayrate contracts tend to become more common for a producer-rig pair as they accumulate experience together, a fact that Corts (2007) attributes to implicit performance incentives. While this empirical regularity is also apparent in the onshore data used in this study, the majority of producer-rig pairs do not change contractual form during the sample period. 40.9% of pairs always use footage contracts, 38.7% always use dayrate contacts, while only 20.5% switch.

The limited contract switching in the data affords an opportunity to verify that the learning results presented thus far are not confounded by changes in contract type. I remove from the sample those producer-rig pairs that switch contracts, and test for relationship-specific learning in the sub-sample of pairs with stable contract types. I include fixed effects for rig-producer interactions in this test to ensure that I do not identify learning effects from cross-pair comparisons, for which contract type may vary. The results of this regression are reported in column V of table 8. I still find a strong and statistically significant effect of joint experience on drilling times. The coefficient on $Experience_{prt}$ is -0.025—very similar to that found when the same regression was run on the full sample (table 7, column III). The joint experience effects I find in the data do not appear to be driven by changes in contract type.[22]

[22] I have run a related regression in which I separate the effect of joint experience into that for producer-rig pairs that use dayrate contracts, and that for pairs that use footage contracts. I find that learning rates are larger for dayrate contracts than for turnkey contracts. However, the endogeneity of contract choice suggests that this is not a causal result. In particular, it may be that learning rates are faster for the types of wells that are amenable to dayrate

Calculation of experience: Finally, I verify that the results are robust to changes in the length of time over which I calculate the experience variables. That is, I ask whether relationship-specific learning is evident when I measure experience using drilling activity over time periods other than two years. The results of these robustness tests are reported in table 9. When I use drilling activity over one year (columns I and II) or three years (columns III and IV), the relationship-specific learning results are very similar to those obtained using the reference case of two years (columns I and IV of table 7). Measuring experience using three years of activity does require that data from 1991 be excluded from the sample, as the rig location data from Smith Bits do not exist prior to 1989. Thus, the sample size in columns III and IV is smaller than that in columns I and II.

Columns V and VI of table 9 report results when experience is measured using five years of drilling activity. These regressions only permit the use of data from 1994 onwards, reducing the sample to 15,731 observations. While the point estimate in column V for $Experience_{prt}$ is larger than that for $Experience_{-prt}$, consistent with relationship-specific learning, the difference between these coefficients is no longer statistically significant (the p-value is 0.364). This decrease in precision from the reference case reflects both the reduction in sample size and the fact that the 5 year experience period is longer than the tenure of the vast majority of rig crew members, as well as some rig foremen.[23] However, when experience is decomposed into field-specific and non-field-specific components (column VI), relationship-specific learning is measured more precisely: pooling of the coefficients on $Experience_{fprt}$ and $Experience_{f\text{-}prt}$ is rejected with a p-value of 0.098.

---

contracting—these wells tend to be more geologically challenging than those drilled under footage contracts and likely present greater scope for learning.

[23] When I run the reference case specification (two-year experience, column I of table 7) using only data from 1994 onwards, the difference between the coefficients on $Experience_{prt}$ and $Experience_{-prt}$ is marginally statistically significant with a p-value of 0.116 (versus 0.036 in the full sample).

7. Empirical Analysis: Relationship Persistence

In this section, I investigate empirically whether the pattern by which producer-rig relationships are formed and broken is consistent with firms' recognition of relationship-specific learning. Specifically, do producers prefer to hire rigs with which they have substantial prior experience? If so, is this preference driven by learning or by other rig-producer specificities?

I execute this analysis using the original SB dataset, prior to its match with the TRRC data. My analysis of relationship persistence does not require the drilling time information in the TRRC data, and the use of the original SB data allows me to examine the complete history of weekly rig locations and relationships from 1991-2005. There are 323,146 rig-week observations in this dataset, and for each I observe the county in which the rig is located and the producer for which the rig is drilling.

Table 10 summarizes the frequency with which rigs either maintain or change their relationships in this sample. Week-to-week, rigs maintain their relationship with their producer 89% of the time. Rigs change producers in 7.5% of the observations, implying that a switch occurs every 13 weeks, on average. Rigs also occasionally exit the market on a temporary or permanent basis; such exits together constitute 3.1% of the data.

I test for relationship persistence by focusing on instances in which a producer has two rigs drilling for it in the same county. I define all such groups of rigs as a "pair," and use these pairs as the unit of observation in my analysis. There are 946 unique pairs in the data, and with two rigs per pair, there exist 1,892 total observations, spread over 844 unique rigs and 553 unique county-producer combinations.[24]

Within each pair, I determine which rig leaves the pair first to work for an alternate producer. I then capture this rig's exit date, and calculate the producer-specific experience of both rigs at that date. I calculate this experience in exactly the same manner as was done for the

[24] I exclude pairs in which both rigs change producers during the same week. I also exclude all pairs in which one or both rigs leave its producer in order to exit the market rather than to work for another firm. This restriction implies that the rig movements I study in my analysis are not driven by a rig's need for maintenance or repairs, or by a rig's failure to win a bid.

relationship-specific learning analysis of section 6. I then test whether the rig that exits first is more likely to be the rig with less relationship-specific experience. This pattern would be consistent with firms' maximization of the benefits of relationship-specific learning.

Figure 5 graphically illustrates the evidence of relationship persistence in this sample. Each point on this plot represents a sub-sample of twenty pairs with similar experience differentials,[25] and the vertical axis indicates the fraction of pairs within each point for which the less experienced rig exited first. The horizontal axis indicates the absolute difference in producer-specific experience (in logs) between the two rigs in each pair, averaged for the twenty pairs in each point. Thus, points plotted on the right side of the plot represent observations in which the two rigs have very different levels of producer-specific experience. For these pairs, it is quite likely that the less-experienced rig will be the first to exit, consistent with firms' recognition of relationship-specific learning's benefits. Meanwhile, observations on the left side of the plot indicate pairs in which the rigs have similar producer-specific experience. In these pairs, the rig with less specific experience appears only slightly more likely to exit first. Thus, when the experience gap between the two rigs in each pair is small, the firms seem relatively indifferent as to which rig exits first.

Figure 5 contrasts with figure 6, in which the horizontal axis measures the difference in *total* drilling experience between the two rigs in each pair rather than the difference in *producer-specific* experience. Here, there appears to be little relationship between the rigs' experience gap and the likelihood that the less-experienced rig exits first. While more experienced rigs are, on average, slightly more likely to be retained than less experienced rigs, this difference is not systematic. This result reflects the fact that the overall experience of a rig does not provide productivity benefits that are producer-specific. While a highly experienced rig may be more

[25] To form the groups of 20 pairs, I first sort the pairs according to the difference in producer-specific experience between the rigs in each pair. I then form groups of 20 from this sorted list, so that each group contains pairs with similar experience differentials.

productive than other rigs, its productivity when working for other producers will also be higher, and it is therefore likely to command a higher price in the market.

Regression analysis confirms these graphical results. I use a conditional logit model to estimate the effect of a rig's producer-specific experience on its probability of being the first to exit its pair. Specifically, I estimate equation (4) below, in which $Experience_{i1}$ denotes the producer-specific experience of rig 1 in pair *i*.

$$\Pr(ExitFirst_{Pair\ i,\ Rig\ 1}) = \frac{\exp(\beta \cdot \log(Experience_{i1}))}{\exp(\beta \cdot \log(Experience_{i1})) + \exp(\beta \cdot \log(Experience_{i2}))} \qquad (4)$$

The results of this regression are reported in column I of table 11: rigs with more producer-specific experience are significantly less likely to exit first. The estimated marginal effect of -0.055 implies that, in a pair consisting of a rig with 12 months of experience and a rig with 1 month of experience, the less experienced rig has a 63.7% probability of being the first to exit. This estimated probability is consistent with figure 5: the difference in the log of experience in this example is 2.48, and the figure suggests that at this difference, the less-experienced rig is likely to exit first with a probability of 60-65%.

Column II of table 11 presents the results of estimating (4) when each rig's total experience is used as the explanatory variable. In this case, there is no significant relationship between experience and movements of rigs between producers. This result is consistent with the scatter of points shown on figure 6. Moreover, column III indicates that the estimated marginal effect of producer-specific experience is not substantially affected when the reference case specification is augmented by including total experience as an additional regressor. In this specification, the coefficient on total experience is positive with a p-value of 0.083.

While the behavior documented in figure 5 and column I of table 11 is consistent with firms' recognition of relationship-specific learning, it could also be consistent with producers' hiring of rigs that are best-suited to drill their wells. That is, if some rigs have producer-specific capabilities due to factors other than learning, then a tendency for producers to work with these rigs could drive the relationship persistence evident in the data. Preferred rigs would be more likely to be both hired first and released last.

Adding fixed effects for rig-producer pairs to (4) could rule out this possibility. Unfortunately, within the 1,892 observations in the sample there are 1,554 unique rig-producer interactions, of which 1,304 occur only once. The limited sample variation remaining after including these fixed effects precludes inference, and it is therefore not possible to explicitly rule out that rig-producer specificities are responsible for the observed relationship persistence.[26]

Nonetheless, there does exist evidence supporting the hypothesis that the relationship persistence is driven by learning effects. First, the analysis of relationship-specific learning in section 6 was robust to the inclusion of rig-producer fixed effects in (2). This robustness suggests either that productivity-improving specificities between a rig and a producer (apart from those generated through learning) are immaterial, or that such specificities do not impact the frequency of rig-producer interactions.

Second, I use information on rig depth ratings to examine directly how the matching of rig attributes to producers' wells affects relationship persistence. Because I observe the depth ratings and drilled well depths for every rig, I can investigate how closely rigs are matched to well depths in practice. Specifically, I return to the dataset of two-rig pairs and calculate the absolute difference between each rig's depth rating and the average depth of the wells it drills within each pair.[27] I then examine whether this "depth difference" affects which rig in each pair changes producers first.

I find that the distribution of the depth difference for the first rig to exit each pair is very similar to that of the rig that exits last: a Kolmogorov-Smirnov test for the equality of these

[26] Estimating (4) with rig-producer fixed effects also creates a severe incidental parameters problem that will cause the estimate of β to be inconsistent (Neyman and Scott 1948, Lancaster 2000). Estimation of a linear probability model with rig-producer fixed effects results in an extremely imprecisely estimated marginal effect: the estimate is +0.062 with a standard error of 0.592.

[27] The average depth difference is 2,313 feet, with a standard deviation of 2,816 feet.

distributions fails to reject equality with a p-value of 0.785. This result indicates that both the first and last rig in each pair to exit is equally well-matched in depth rating to its wells.

I confirm this conclusion by augmenting regression (4) with each rig's depth difference. As shown in column IV of table 11, the estimated coefficient on depth difference is very close to zero and statistically insignificant. More importantly, the primary finding that rigs with relatively high producer-specific experience are less likely to change producers is unaffected. This result indicates that the relationship persistence in the data is not driven by producers choosing to work more frequently with rigs that are better sized for their wells. Instead, the accumulated evidence suggests that producers' tendency to employ rigs with which they have significant prior experience is designed to maximize the benefits of relationship-specific learning.

8. Conclusions

The empirical results presented in this paper identify relationship-specific learning as an important driver of productivity gains in the oil and gas drilling industry. While the fact that such an effect exists is not in itself particularly surprising—drilling industry participants have told me that they believe relationship-specific learning occurs—its importance has not received attention in the learning-by-doing literature. The primary contribution of this paper is the establishment of both the significant magnitude of this learning and its broader economic relevance. I find that a rig that accumulates experience with one producer improves its productivity twice as quickly as a rig that frequently changes contracting partners. This large productivity benefit gives producers and rigs a strong incentive to maintain their relationships. Accordingly, the data demonstrate that producers are more likely to work with rigs with which they have substantial prior experience than those with which they have worked relatively little.

These results seem likely to generalize to other industries in which outsourcing is common. For example, construction contractors or management consulting firms may develop relationship-specific intellectual capital through joint work experience with their clients. The importance of relationship-specific learning presumably varies with industry and firm characteristics. For example, greater technical complexity in an industry's production process could drive steeper learning curves than those documented in this paper. Firms may also be able to take actions that influence their rate of relationship-specific learning. A lead firm might embed some of its employees within the organizations of its contracting partners, or a contractor might set up offices near its clients.

While such actions plausibly increase the rate of learning, the accumulated knowledge that results is also a form of relationship-specific capital. As such, it may amplify opportunities for *ex post* rent-seeking. It is therefore in the interest of firms to develop contracting arrangements that alleviate this problem when joint experience effects are significant. In the extreme, firms may need to integrate to fully capture the benefits of relationship-specific learning. This potential efficiency benefit of integration may merit consideration in merger analysis.

Finally, I find that horizontal learning spillovers are unlikely to be important in oil and gas drilling. Given prior findings of spillover effects in semiconductor manufacturing and shipbuilding, this result suggests that the importance of spillovers varies with industry characteristics. For example, the lack of spillovers in oil and gas drilling may be related to the competitive nature of production from a common pool resource. Because economic theory indicates that learning spillovers are important for macroeconomic growth, obtaining a deeper understanding of spillovers' determinants is an objective worthy of further research.

Appendix: Forgetting Effects

In this appendix, I investigate whether experience effects in the drilling industry decay over time, consistent with institutional "forgetting" of knowledge. That is, I ask whether experience from the distant past has a smaller effect on current productivity than does recent experience. This inquiry relates to research by Argote *et al.* (1990) and Benkard (2000) that identified forgetting effects in shipbuilding and aircraft manufacturing, respectively.

Throughout the text of this paper, I define $Experience_{fpt}$—the experience of producer p in field f at date t—as the number of wells drilled by producer p in field f during the two years prior to t. Here, I instead define $Experience_{fpt}$ as a function of a decay parameter α, per expression (A1) below.

$$Experience_{fpt}(\alpha) = \sum_{\tau=t-730}^{t} 1\{Well\ Completed\ at\ \tau\} e^{-\alpha(t-\tau)/365} \qquad \text{(A1)}$$

For positive values of α, wells drilled on dates long before date t carry less weight in $Experience_{fpt}(\alpha)$ than do wells drilled near date t. I estimate α by inserting (A1) into learning specification (1), yielding the nonlinear expression (A2) below.

$$log(DrillTime_{fpt}) = \beta \cdot log(Experience_{fpt}(\alpha)) + \gamma_f + \delta_p + \nu_t + \boldsymbol{\varphi X_{fpt}} + \varepsilon_{fpt} \qquad \text{(A2)}$$

I estimate (A2) using a method of moments estimator, in which I allow ε_{fpt} to be heteroskedastic and correlated within each field. As in the text, I instrument for experience using the number of days producer p actively drilled in field f during the two years prior to date t.[28]

I obtain a point estimate of α equal to 1.823 with a clustered standard error of 0.862, suggesting the presence of forgetting. While this result suggests, at first glance, that $\hat{\alpha}$ is statistically significant at the 5% level, the proper statistical test for forgetting effects must take into account the negative correlation between $\hat{\alpha}$ and $\hat{\beta}$. The point estimate of β is -0.055,

[28] My estimation also instruments for the derivative of experience with respect to α using the derivative of drilling time-based experience with respect to α, though this does not significantly affect the results.

substantially larger in magnitude than in the reference case results (table 4, column I). This increase in magnitude occurs because learning is now a function of depreciated experience rather than total experience. Because the estimates of α and β are linked in this way, I test for forgetting by testing whether these estimates are jointly different from those reported in the reference case, in which $\alpha = 0$ and $\beta = -0.036$. I find that the null hypothesis of no forgetting can be rejected at only the 10% level (the p-value is 0.091) rather than at the 5% level.

The estimated rate of experience depreciation is somewhat large: the point estimate of α implies that a well drilled one year ago makes a contribution to experience that is only 16% of that made by a well drilled one day ago. This depreciation rate is not as great as that estimated by Argote *et al*. (1990) in shipbuilding (for which the corresponding figure is 3.2%), though greater than that estimated by Benkard (2000) in aircraft manufacturing (61%). While this result could reflect literal human forgetting of knowledge or turnover amongst producers' personnel, it may also reflect losses of intellectual capital associated with changes in producers' drilling rigs. That is, the forgetting effect estimated in this appendix could be driven by the relationship-specificity of learning in the drilling industry. I investigate this possibility by augmenting (A2) with rig fixed effects and variables measuring rig and relationship-specific experience. I find that controlling for relationship-specific capital obscures producers' forgetting because changes in this capital are highly collinear with experience depreciation. While the new point estimate of α is still positive and quite large at 7.426, it is estimated extremely imprecisely: the standard error is 13.37. This collinearity suggests that losses of relationship-specific capital between lead firms and contractors may be a mechanism by which experience effects are forgotten.

References

Adeleye, A.R., B.K. Virginillo, A.W. Iyoho, K. Parenteau, and Henry Licis, "Improving Drilling Performance Through Systematic Analysis of Historical Data: Case Study of a Canadian Field," Society of Petroleum Engineers paper 87177 (Mar., 2004).

Alchian, Armen, "Reliability of Progress Curves in Airframe Production," *Econometrica* 31 (Oct., 1963), 679-694.

Arellano, Manuel, "Computing Robust Standard Errors for Within-Groups Estimators," *Oxford Bulletin of Economics and Statistics* 49 (Nov., 1987), 431-434.

Argote, Linda, Sara L. Beckman, and Dennis Epple, "The Persistence and Transfer of Learning in Industrial Settings," *Management Science* 36 (Feb., 1990), 140-154.

Benkard, C. Lanier, "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review* 90 (Sep., 2000), 1034-1054.

Brett, J.F. and K.K. Millheim, "The Drilling Performance Curve: A Yardstick for Judging Drilling Performance," Society of Petroleum Engineers paper 15362 (Oct., 1986).

Coase, Ronald H. "The Nature of the Firm," *Economica* 4 (Nov., 1937), 386-405.

Corts, Kenneth S. "The Interaction of Implicit and Explicit Contracts in Construction and Procurement Contracting," University of Toronto working paper (Jun., 2007).

Corts, Kenneth S. and Jasjit Singh. "The Effect of Repeated Interaction on Contract Choice: Evidence from Offshore Drilling," *Journal of Law, Economics, and Organization* 20 (2004), 230-260.

Gibbons, Robert, "Four Formal(izable) Theories of the Firm?" *Journal of Economic Behavior and Organization* 58 (2005), 200-245.

Groysberg, Boris, "Can They Take it With Them? The Portability of Star Knowledge Workers' Performance: Myth or Reality?" Harvard University working paper (Dec., 2001).

Hendricks, Kenneth and Robert H. Porter, "An Empirical Study of an Auction with Asymmetric Information," *American Economic Review* 78 (Dec., 1988), 865-883.

Huckman, Robert S. and Gary P. Pisano, "The Firm Specificity of Individual Performance: Evidence from Cardiac Surgery," *Management Science* 52 (Apr., 2006), 473-488.

Irwin, Douglas A. and Peter J. Klenow, "Learning-by-Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy* 102 (Dec., 1994), 1200-1227.

Joskow, Paul L. and Nancy L. Rose, "The Effects of Technological Change, Experience, and Environmental Regulation on the Construction Cost of Coal-Burning Generating Units," *RAND Journal of Economics* 16 (Spring, 1985), 1-27.

Jovanovic, Boyan and Yaw Nyarko, "Learning by Doing and the Choice of Technology," *Econometrica* 64 (Nov. 1996), 1299-1310.

Klein, Benjamin, Robert G. Crawford, and Armen A. Alchian, "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process," *Journal of Law and Economics* 21 (Oct., 1978), 297-326.

Lafontaine, Francine and Margaret Slade, "Vertical Integration and Firm Boundaries: The Evidence," *Journal of Economic Literature* 45 (Sep., 2007), 629-685.

Lancaster, Tony, "The Incidental Parameter Problem Since 1948," *Journal of Econometrics* 95 (2000), 391-413.

Managing Automation, "Boeing's Big Supply Chain Wager" (17 May, 2007).

Neyman, Jerzy and Elizabeth L. Scott, "Consistent Estimates Based on Partially Consistent Observations," *Econometrica* 16 (Jan., 1948), 1-32.

Parente, Stephen L., "Technology Adoption, Learning-by-Doing, and Economic Growth," *Journal of Economic Theory* 63 (Aug., 1994), 346-369.

Thornton, Rebecca A. and Peter Thompson, "Learning from Experience and Learning from Others: An Exploration of Learning and Spillovers in Wartime Shipbuilding," *American Economic Review* 91 (Dec., 2001), 1350-1368.

Whinston, Michael D., "On the Transaction Cost Determinants of Vertical Integration," *Journal of Law, Economics, and Organization* 19 (2003), 1-23.

Wiggins, Steven N. and Gary D. Libecap, "Oil Field Unitization: Contractual Failure in the Presence of Imperfect Information," *American Economic Review* 75 (Jun., 1985), 368-385.

Williamson, Oliver E., *Markets and Hierarchies: Analysis and Antitrust Implications,* New York: Free Press (1975).

Williamson, Oliver E., "Transaction-Cost Economics: The Governance of Contractual Relations," *Journal of Law and Economics* 22 (Oct., 1979), 233-261.

Williamson, Oliver E., *The Economic Institutions of Capitalism,* New York: Free Press (1985).

Wooldridge, Jeffrey M., "Cluster-Sample Methods in Applied Econometrics," *American Economic Review Papers and Proceedings* 93 (May, 2003), 133-138.

Wright, T.P., "Factors Affecting the Cost of Airplanes," *Journal of the Aeronautical Sciences* 3 (Feb., 1936), 122-128.

Figure 1: Photo of drilling rig

Table 1: Distributions of wells, fields, producers, and rigs

	Min	25th percentile	Median	75th percentile	Mean	Max
Number of wells per field	2	2	4	8	11.9	715
Number of wells per producer	2	3	7	20	24.6	698
Number of wells per rig	2	4	8	19	14.4	161
Number of producers per field	1	1	2	3	2.8	60
Number of fields per producer	1	2	3	6	5.8	111
Number of rigs per driller	1	1	3	6	7.7	195

Figure 2: Drilling times vs. well depths

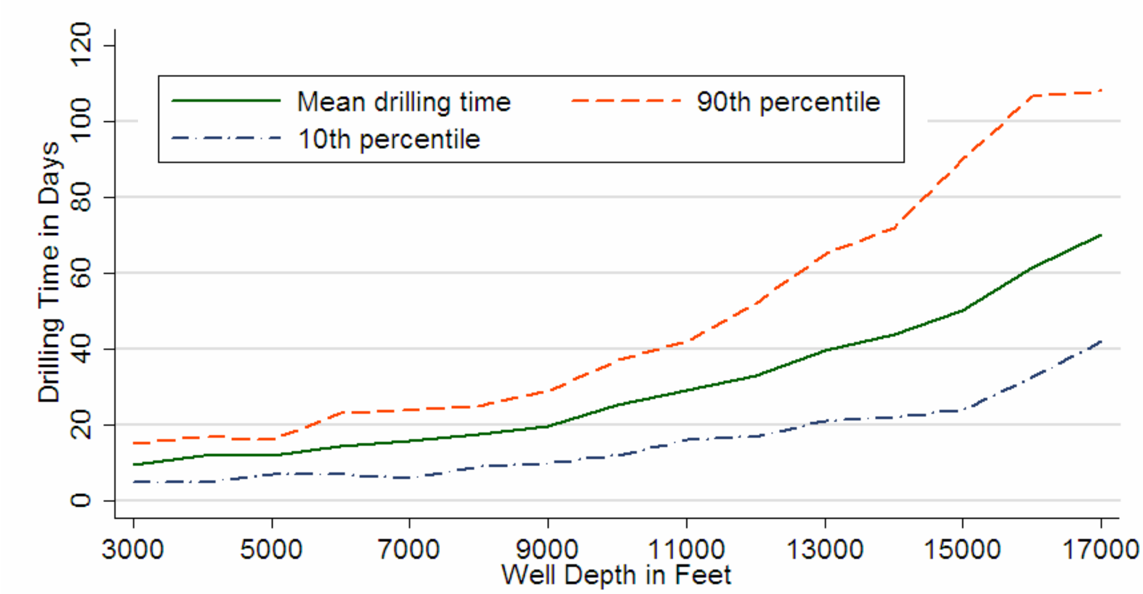

Observations are grouped into depth "bins" of 1000 feet.
Horizontal axis excludes highest and lowest 1% of depths.

Table 2: Sample summary statistics

	Number of observations	Min	Median	Mean	Std. Dev.	Max
Drilling time (days)	19714	2	18	23.7	20.8	180
Well depth (feet)	19174	631	9000	9040	2829	23000
Oil well (0/1 dummy)	19174	0	0	0.381	0.486	1
Gas well (0/1 dummy)	19174	0	1	0.600	0.490	1
Oil and gas well (0/1 dummy)	19174	0	0	0.020	0.138	1
Dry hole (0/1 dummy)	19174	0	0	0.082	0.274	1

Table 3: Summary statistics of field and producer experience variables

Number of wells drilled during the past two years in:	Number of observations	Min	Median	Mean	Std. Dev.	Max
Same field, same producer	19714	1	6	17.6	40.6	552
Different field, same producer	19174	1	40	115.3	156.2	918
Same field, different producer	19174	1	9	66.7	169.8	797

The well represented by each observation is included in all measures of experience. Thus, the minimum experience level is one rather than zero.

Figure 3: Estimates of learning by field producers: spline and log-log specifications

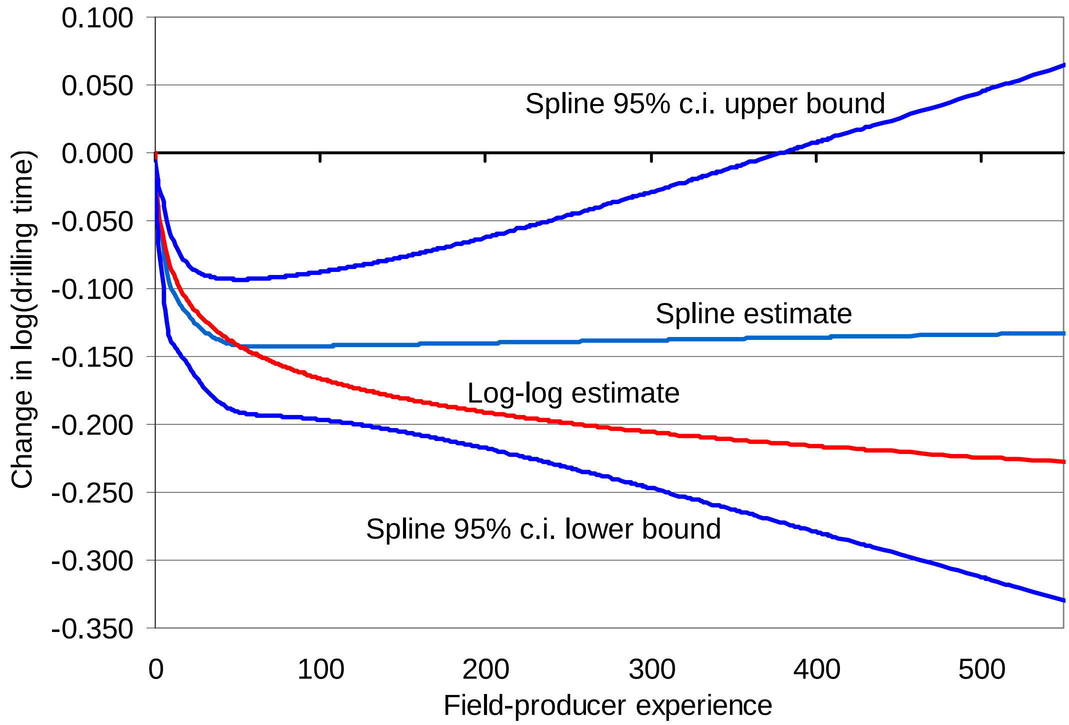

Figure 4: Estimates of learning by field producers: spline and log-log specifications; wells with fewer than 100 units of experience

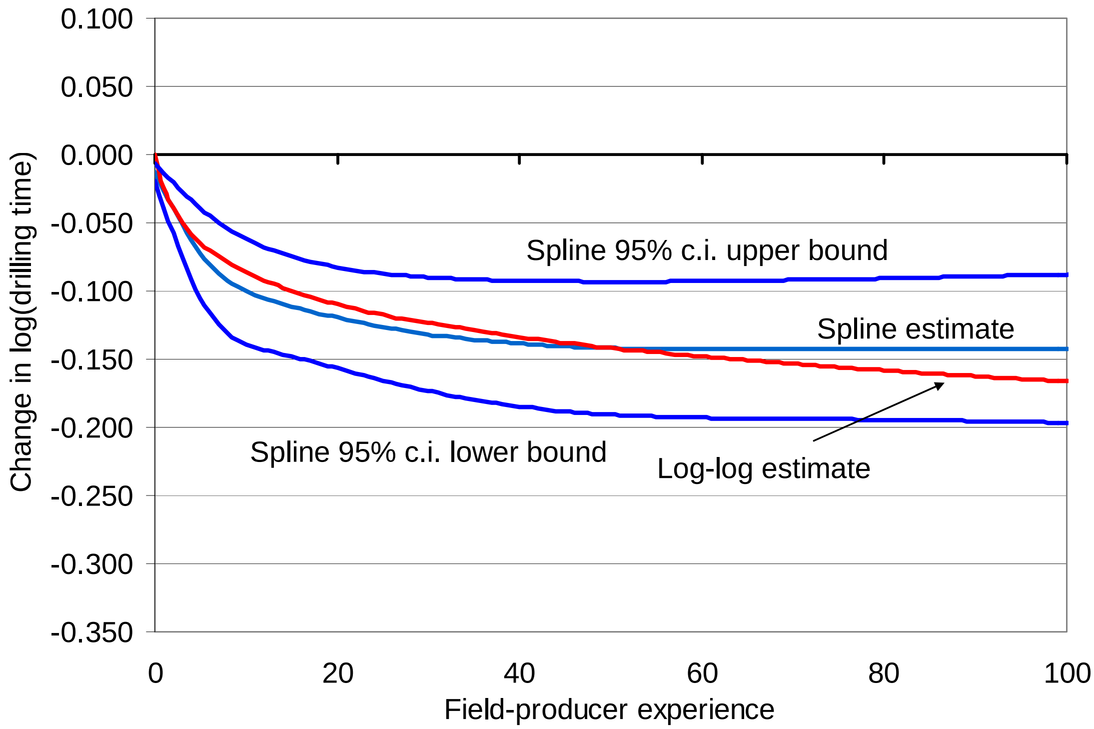

Table 4: Regression results for learning by field producers
Dependent variable is log(drilling time)

	Point estimate	Standard error
log(field-producer experience)	-0.036	(0.005) ***
log(well depth)	1.181	(0.077) ***
Gas well dummy	0.041	(0.025)
Oil and gas well dummy	0.034	(0.051)
Dry hole dummy	0.144	(0.027) ***
February dummy	-0.016	(0.017)
March dummy	-0.026	(0.015) *
April dummy	-0.036	(0.016) **
May dummy	-0.037	(0.016) **
June dummy	-0.053	(0.016) ***
July dummy	-0.030	(0.016) *
August dummy	-0.029	(0.015) *
September dummy	0.005	(0.016)
October dummy	-0.033	(0.016) **
November dummy	-0.032	(0.016) **
December dummy	-0.019	(0.016)
1992 dummy	-0.041	(0.025)
1993 dummy	-0.003	(0.025)
1994 dummy	-0.034	(0.031)
1995 dummy	-0.033	(0.035)
1996 dummy	-0.023	(0.033)
1997 dummy	-0.008	(0.036)
1998 dummy	-0.016	(0.040)
1999 dummy	-0.082	(0.039) **
2000 dummy	-0.074	(0.038) *
2001 dummy	-0.004	(0.041)
2002 dummy	-0.059	(0.040)
2003 dummy	-0.099	(0.043) **
2004 dummy	-0.159	(0.045) ***
2005 dummy	-0.167	(0.046) ***

Regression includes fixed effects for fields and producers
Standard errors are clustered on field.
*,**,*** indicate significance at the 10%, 5%, and 1% level.

Table 5: Regression results for learning by field producers
Dependent variable is log(drilling time)

Log of experience with:	I Reference Case (Table 4)	II Experience not instrumented	III Field-producer fixed effects	IV Learning spillovers
Same field, same producer	-0.036***	-0.047***	-0.047***	-0.036***
	(0.005)	(0.005)	(0.007)	(0.005)
Different field, same producer	-	-	-	-0.013*
	-	-	-	(0.007)
Same field, different producer	-	-	-	-0.006
	-	-	-	(0.006)
Controls for depth and well type	Y	Y	Y	Y
Month and year dummies	Y	Y	Y	Y
Field and producer dummies	Y	Y	Y	Y
Field X producer dummies	N	N	Y	N
Number of observations	19174	19174	19174	19174

Parenthetical values indicate standard errors clustered on field.
*,**,*** indicate significance at the 10%, 5%, and 1% level.
Estimated coefficients on depth and well type variables vary little across the specifications.

Table 6: Summary statistics of rig experience variables

Number of weeks of drilling within past two years by:	Number of observations	Min	Median	Mean	Std. Dev.	Max
Same rig, different producer	19174	1	34	37.1	28.9	105
Same rig, same producer	19174	1	12	24.4	28.6	105
Same rig, diff producer, diff field	19174	1	22	28.5	26.4	102
Same rig, same producer, diff field	19174	1	1	11.2	19.3	103
Same rig, diff producer, same field	19174	1	6	9.7	12	105
Same rig, same producer, same field	19174	1	8	15.2	19.1	105

The well represented by each observation is included in all measures of experience. Thus, the minimum experience level is one rather than zero.

Table 7: Regression results for relationship-specific learning
Dependent variable is log(drilling time)

Log of experience with:	I Reference case	II Drilling company learning	III Rig-producer F.E.	IV Rig-field specificities	V Rig-producer-field F.E.
Same field, same producer	-0.022***	-0.021***	-0.019**	-0.013**	-0.008
($Experience_{fpt}$)	(0.006)	(0.006)	(0.008)	(0.006)	(0.017)
Different field, same producer	-0.010	-0.011	0.010	-0.013	0.001
($Experience_{-fpt}$)	(0.008)	(0.008)	(0.012)	(0.008)	(0.016)
Same field, different producer	-0.006	-0.006	0.003	-0.005	0.014
($Experience_{f-pt}$)	(0.006)	(0.006)	(0.009)	(0.006)	(0.015)
Same rig, different producer	-0.010**	-0.012***	0.005	-	-
($Experience_{-prt}$)	(0.004)	(0.004)	(0.006)	-	-
Same rig, same producer	-0.021***	-0.022***	-0.024***	-	-
($Experience_{prt}$)	(0.004)	(0.005)	(0.005)	-	-
Same drilling company, different rig	-	0.019*	-	-	-
	-	(0.010)	-	-	-
Same drilling company, same producer	-	0.001	-	-	-
	-	(0.005)	-	-	-
Same rig, diff field, diff producer	-	-	-	-0.007*	9.5E-05
($Experience_{-f-prt}$)	-	-	-	(0.004)	(0.007)
Same rig, same field, diff producer	-	-	-	-0.007	0.001
($Experience_{f-prt}$)	-	-	-	(0.006)	(0.009)
Same rig, diff field, same producer	-	-	-	-0.002	-0.004
($Experience_{-fprt}$)	-	-	-	(0.004)	(0.007)
Same rig, same field, same producer	-	-	-	-0.036***	-0.040***
($Experience_{fprt}$)	-	-	-	(0.005)	(0.011)
Rig X producer dummies	N	N	Y	N	N
Rig X producer X field dummies	N	N	N	N	Y
Number of observations	19174	19174	19174	19174	19174

Parenthetical values indicate standard errors clustered on field.

*,**,*** indicate significance at the 10%, 5%, and 1% level.

All regressions include controls for depth and well type, as well as field, producer, and rig fixed effects.

Table 8: Robustness tests for relationship-specific learning results
Dependent variable is log(drilling time)

Log of experience with:	I Wells 8000 ft and deeper	II Wells 8000 ft and deeper	III Field-producer spline	IV Field-producer spline	V Stable contract type
Same field, same producer	-0.017***	-0.012	spline	spline	-0.022**
(*Experience* $_{fpt}$)	(0.007)	(0.008)			(0.010)
Different field, same producer	-0.012	-0.013	-0.010	-0.014	0.019
(*Experience* $_{-fpt}$)	(0.009)	(0.009)	(0.008)	(0.008)	(0.014)
Same field, different producer	-0.007	-0.007	-0.006	-0.005	-3.4E-04
(*Experience* $_{f\text{-}pt}$)	(0.008)	(0.008)	(0.006)	(0.006)	(0.010)
Same rig, different producer	-0.010**	-	-0.010**	-	0.003
(*Experience* $_{-prt}$)	(0.004)	-	(0.004)	-	(0.007)
Same rig, same producer	-0.020***	-	-0.021***	-	-0.025***
(*Experience* $_{prt}$)	(0.005)	-	(0.004)	-	(0.007)
Same rig, diff field, diff producer	-	-0.011**	-	-0.007*	-
(*Experience* $_{-f\text{-}prt}$)	-	(0.004)	-	(0.004)	-
Same rig, same field, diff producer	-	0.001	-	-0.007	-
(*Experience* $_{f\text{-}prt}$)	-	(0.007)	-	(0.006)	-
Same rig, diff field, same producer	-	-0.005	-	-0.002	-
(*Experience* $_{-fprt}$)	-	(0.004)	-	(0.004)	-
Same rig, same field, same producer	-	-0.029***	-	-0.037***	-
(*Experience* $_{fprt}$)		(0.006)		(0.005)	
Rig X producer dummies	N	N	N	N	Y
Number of observations	12654	12654	19174	19174	15276

Parenthetical values indicate standard errors clustered on field.

*,**,*** indicate significance at the 10%, 5%, and 1% level

All regressions include controls for depth and well type, as well as field, producer, and rig fixed effects

Table 9: Robustness tests for relationship-specific learning results: changes to calculation of experience
Dependent variable is log(drilling time)

Log of experience with:	I One-year experience	II One-year experience	III Three-year experience	IV Three-year experience	V Five-year experience	VI Five-year experience
Same field, same producer	-0.030***	-0.025***	-0.017***	-0.010	-0.017***	-0.010
(*Experience* $_{fpt}$)	(0.006)	(0.007)	(0.004)	(0.006)	(0.006)	(0.007)
Different field, same producer	-0.010	-0.003	-0.009	-0.012	0.002	-0.005
(*Experience* $_{-fpt}$)	(0.008)	(0.006)	(0.008)	(0.008)	(0.009)	(0.010)
Same field, different producer	-0.003	-0.011	-0.005	-0.005	-0.001	0.003
(*Experience* $_{f-pt}$)	(0.006)	(0.008)	(0.007)	(0.007)	(0.009)	(0.009)
Same rig, different producer	-0.011***	-	-0.010***	-	-0.015***	-
(*Experience* $_{-prt}$)	(0.004)	-	(0.005)	-	(0.006)	-
Same rig, same producer	-0.020***	-	-0.020***	-	-0.021***	-
(*Experience* $_{prt}$)	(0.004)	-	(0.004)	-	(0.004)	-
Same rig, diff field, diff producer	-	-0.012***	-	-0.008**	-	-0.007
(*Experience* $_{-f-prt}$)	-	(0.004)	-	(0.003)	-	(0.004)
Same rig, same field, diff producer	-	-0.002	-	-0.005	-	-0.012
(*Experience* $_{f-prt}$)	-	(0.005)	-	(0.007)	-	(0.008)
Same rig, diff field, same producer	-	-0.008	-	-0.002	-	-0.003
(*Experience* $_{-fprt}$)	-	(0.004)	-	(0.004)	-	(0.004)
Same rig, same field, same producer	-	-0.029***	-	-0.028***	-	-0.029***
(*Experience* $_{fprt}$)		(0.006)		(0.005)		(0.005)
Number of observations	19174	19174	17987	17987	15731	15731

Parenthetical values indicate standard errors clustered on field.

*,**,*** indicate significance at the 10%, 5%, and 1% level

All regressions include controls for depth and well type, as well as field, producer, and rig fixed effects

Table 10: Summary of rig movements

	Number of observations (rig-weeks)	Percent of total
Continue working for same producer	288,801	89.4%
Change producers	24,395	7.5%
Exit from market for at least two months	7,087	2.2%
Exit from market permanently	2,863	0.9%
Total	323,146	100.0%

Figure 5: Likelihood that the least experienced rig is the first to change producers vs. the within-group difference in rigs' producer-specific experience

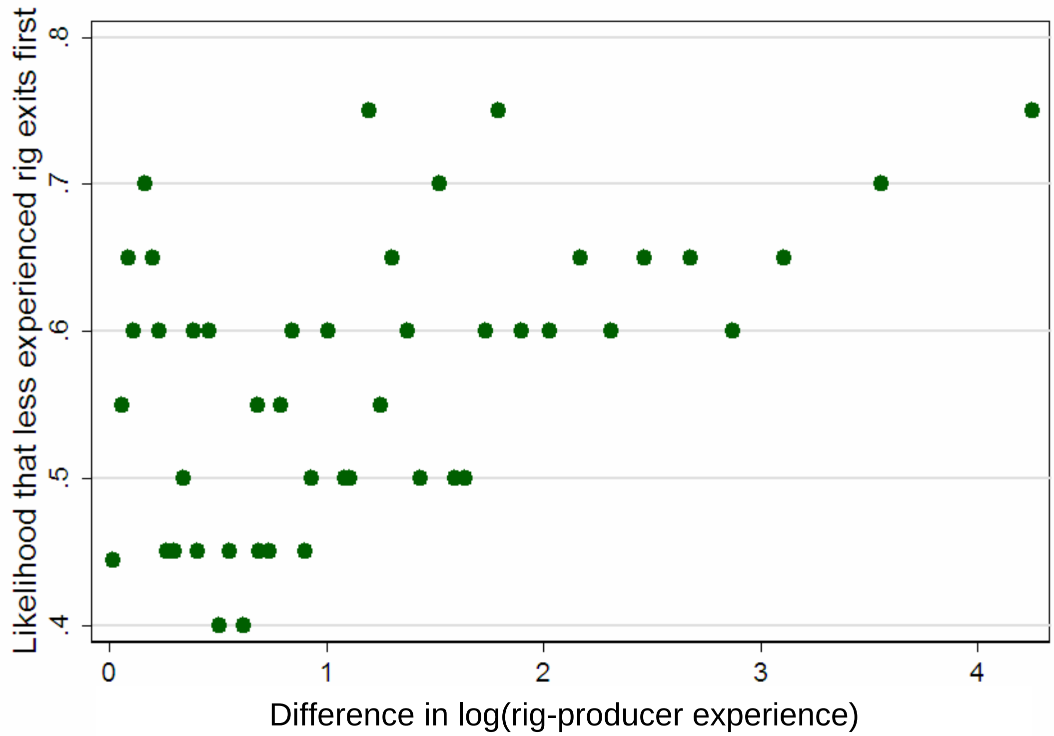

Figure 6: Likelihood that the least experienced rig is the first to change producers vs. the within-group difference in rigs' overall experience

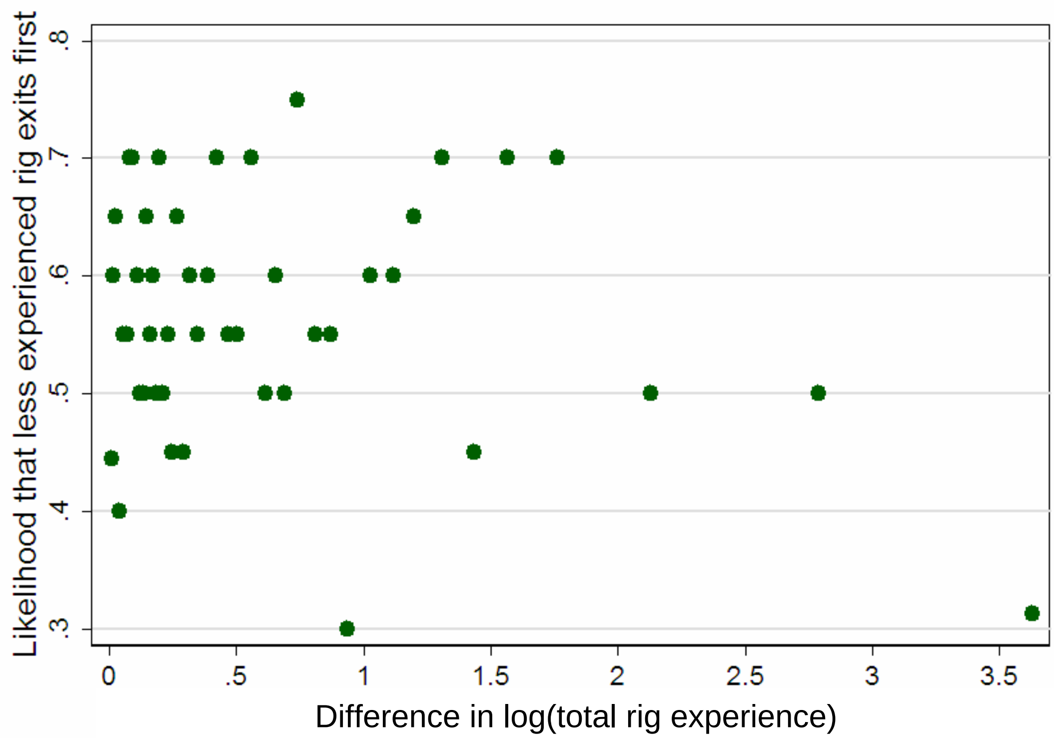

Table 11: Conditional logit estimates for the probability a rig is the first to exit its pair
Values shown are marginal effects: dPr(ExitFirst) / dX

	I	II	III	IV
	Reference case estimate	Total rig experience	Producer-specific and total experience	Control for depth rating
Log of rig's producer-specific experience	-0.055***	-	-0.062***	-0.055***
	(0.013)	-	(0.014)	(0.013)
Log of rig's total experience	-	-0.003	0.030*	-
	-	(0.016)	(0.017)	-
Average absolute difference between rig depth rating and well depths ('000 feet)	-	-	-	0.001
	-	-	-	(0.005)
Number of rig pairs	946	946	946	946

Marginal effects calculated at sample mean.

Parenthetical values indicate standard errors clustered on producer.

*,**,*** indicate significance at the 10%, 5%, and 1% level.