

NBER WORKING PAPER SERIES

LEARNING BY DRILLING:
INTER-FIRM LEARNING AND RELATIONSHIP PERSISTENCE IN THE TEXAS OILPATCH

Ryan Kellogg

Working Paper 15060
<http://www.nber.org/papers/w15060>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2009

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, Jennifer Brown, Meghan Busse, Michael Greenstone, Jun Ishii, Kai-Uwe Kühn, Jeffrey Perloff, and numerous seminar participants. Tay Feder provided excellent research assistance. I am particularly grateful to Richard Mason of RigData and Turner Williamson of the Texas Railroad Commission for many illuminating conversations about the onshore Texas drilling industry. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Ryan Kellogg. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch
Ryan Kellogg
NBER Working Paper No. 15060
June 2009
JEL No. D24,L14,L71,Q40

ABSTRACT

This paper examines the importance of learning-by-doing that is specific not just to individual firms, but to pairs of firms working together in a contracting relationship. Using new, detailed data from the oil and gas industry, I find that the joint productivity of an oil production company and its drilling contractor is enhanced significantly as they accumulate experience working together. This learning is relationship-specific: drilling rigs generally cannot fully appropriate the productivity gains acquired through experience with one production company to their work for another. This result is robust to other ex ante match specificities.

Relationship-specific learning is consequential because it implies that relationship stability is important to productivity. When two firms accumulate experience working together, relationship-specific intellectual capital is created that cannot be appropriated to pairings with other firms. If the relationship is broken, this capital is destroyed and productivity decreases, thereby giving firms an incentive to maintain long-term relationships. Accordingly, the data indicate that production companies prefer to work with drilling rigs which they have accumulated considerable experience rather than those with which they have worked relatively little. I demonstrate that this contracting pattern is difficult to explain with switching costs or ex ante match specificities alone.

Ryan Kellogg
University of Michigan
Department of Economics
238 Lorch Hall
611 Tappan Street
Ann Arbor, MI 48109-1220
and NBER
kelloggr@umich.edu

1. Introduction

The economic consequences of learning-by-doing, the hypothesis that unit costs decrease with cumulative production, are well-known. In industrial organization, learning-by-doing can rationalize pricing below short-run marginal cost and lead to increases in industry concentration through the emergence of a low-cost dominant firm (Spence 1981, Cabral and Riordan 1994, Benkard 2004, Besanko *et al.* 2007). In macroeconomics, on-the-job learning and knowledge spillovers are widely believed to play important roles in driving endogenous economic growth (Arrow 1962, Lucas 1988, Stokey 1988, Parente 1994, Jovanovic and Nyarko 1996). This paper uses a new dataset to document a form of learning that has thus far received little attention: relationship-specific learning. Learning-by-doing can be relationship-specific when the productivity improvements associated with the accumulation of experience are specific to not just an individual firm but to pairs of firms working together in a contracting relationship. For example, a contract 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.

Relationship-specific learning is consequential because it implies that relationship stability is important to productivity. When two firms accumulate experience working together, relationship-specific intellectual capital is created that cannot be appropriated to pairings with other firms. If the relationship is broken, this capital is destroyed and productivity decreases. Relationship-specific learning therefore gives firms an incentive to work with contractors with which they have substantial experience rather than those with which they have worked relatively little. This learning may also be a mechanism behind recent documentation of forgetting effects. Argote, Beckman, and Epple (1990), Benkard (2000), and Thompson (2003) find evidence that a firm's recent production experience has a stronger impact on productivity than does older experience. Some of this experience depreciation may reflect an unobserved change in the firm's contracting relationships and a commensurate loss of relationship-specific capital. Finally, like forgetting, relationship-specific learning may be important at the macroeconomic level: recessions that disrupt production and fracture relationships may result in a productivity decrease that persists beyond the rebound in output during the recovery.

Are relationship-specific learning effects sufficiently large that they plausibly play a role in firms' contracting or are a determinant of economic productivity? The literature is largely silent on this question. This paper therefore empirically evaluates the importance of relationship-specific learning using a new dataset from the U.S. onshore oil and gas drilling industry. I ask two questions. First, when production requires coordinated inputs from multiple firms, to what

extent is productivity a function of not just each firm's individual experience but also the firms' joint experience? And second, do firms prefer to maintain long-term relationships rather than regularly switch contracting partners, consistent with a desire to maximize relationship-specific learning's productivity benefits?

The U.S. onshore drilling industry is well-suited to this investigation for several reasons. First, drilling requires inputs from two types of firms: production companies ("producers") and drilling companies. 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 requires the technical optimization of drilling procedures as well as teamwork between producer personnel and the rig crew—skills that may be acquired through experience. Third, I have collected excellent data on both drilling contracting and performance, covering nearly 20,000 wells drilled over 1991-2005, with which I can 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 more than twice as large as those of a rig that frequently changes producers. Because I observe multiple wells drilled per producer-rig pair, I am able to distinguish this learning effect from any *ex ante* match specificities that might cause certain firm pairs to drill more effectively and more frequently than others.

For the average well in my dataset, I estimate that relationship-specific learning improves drilling productivity by 3.8%, yielding cost savings of about \$9,700 per well. These savings give firms an incentive to maintain long-term relationships. Accordingly, the data indicate that producers prefer to work with rigs which they have accumulated considerable experience rather than those with which they have worked relatively little. I demonstrate that this contracting pattern is difficult to explain with switching costs or *ex ante* match specificities alone.

Beyond these primary results, I also test for the presence of learning spillovers in the drilling industry. While I find that oil and gas producers are able to transfer knowledge obtained from experience in one oilfield to the drilling of another, I do not find evidence supporting learning spillovers across producers working side-by-side in the same field. This result stands in contrast to other studies that identify modest cross-firm learning spillovers in semiconductor manufacturing and shipbuilding (Irwin and Klenow 1994, Thornton and Thompson 2001).

While this paper focuses on the oil and gas drilling industry, it seems likely that the prevalence of relationship-specific learning extends well beyond the oilpatch. Construction and large manufacturing projects, for example, regularly involve multiple contractors and sub-contractors working under a lead, general contractor. 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 commented that “interpersonal communication skills and building relationships have become more important than ever” (*Managing Automation* 2007).

Finally, while I focus on relationship-specific learning as a phenomenon that occurs between firms, learning specificities are likely to be important within firms as well: workers may develop skills that are specific to their particular employer. For example, Huckman and Pisano (2006) find evidence suggesting that doctors' surgical outcomes depend more on their hospital-specific experience than on their general experience. Becker (1964), Prendergast (1993), and Gibbons and Waldman (1999, 2004) discuss the implications of job-specific learning for equilibrium wage and promotion paths, explaining why, for example, wages increase with age at a decreasing rate.¹ These theories could in principle be translated to the pricing of service contracts between firms when learning is relationship-specific.

The remainder of the paper is organized as follows: section 2 provides background information on the oil and gas drilling industry, and section 3 discusses industry mechanisms 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 analysis 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.

¹ For an example and survey of the empirical literature on wage dynamics, see Poletaev and Robinson (2008).

2. The Onshore Oil and Gas Drilling Industry

2.1 Production companies and the drilling problem

Oil and gas reserves are found in geologic formations known as fields that lie beneath the earth's surface. 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.² 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.³

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. Once a well is drilled to its target depth, the drilling rig is no longer needed and the well, if successful, will produce oil and/or gas for a period of several years.

Typically, there is significant variation in geology across fields, particularly with regards to the depth at which they are buried. Some fields are found as shallow as 3,000 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 types of rock that must be drilled through also vary considerably: the layers of sandstone, shale, and limestone that may be encountered in one area will generally not be the same as those found elsewhere.

Wells fall into two broad categories. "Wildcats" are those 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 oil and gas production.

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

² Onshore leasing differs from the federal offshore leasing studied by Hendricks and Porter (1988) in that there is no centralized process in Texas by which producers obtain leases. Instead, producers proactively approach the holders of mineral rights, who may then negotiate lease terms or organize a competitive bidding process.

³ 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.

knowledge from seismic imaging and previously drilled wells, and are therefore better placed to make these engineering decisions.⁴

2.2 Rigs and contracting

The actual drilling of wells is conducted 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 equipment that eases the drilling of horizontal and directional wells.

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. My interviews with 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 temporal variation with which producers drill wells. The drilling activity of any producer fluctuates with oil production outcomes from recently drilled wells and the firm's success in finding new fields. Successful wildcats and development wells often lead to additional drilling, while "dry" unproductive holes do not. The mobility and non-specificity of rigs allow them to smooth these fluctuations in drilling requirements across 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 since they are generally reluctant to commit to a long-term contract when the total number of wells they will drill is unknown and contingent on oil production from the first several wells drilled. For

⁴ 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.

example, if a drilled well turns out to be a dry hole, the producer will usually not want to follow-up with additional drilling in the field. Long-term relationships are therefore generally maintained through repeat contracting rather than formal long-term contracts.⁵ To the extent that relationship-specific learning is important, this repeat contracting creates rents that can be bargained over at each renewal. However, unlike classic examples of relationship-specific investments from the transactions cost literature (Williamson 1975, 1985 and Klein *et al.* 1978), relationship-specific learning does not generally require costly up-front investment so that the lack of a long-term contract does not create an inefficiency.⁶

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 for instance the well's total depth 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. In cases where a producer is following-up an initial well with further drilling and wishes to retain its current rig, it will generally renew its current contract rather than hold another auction.

The RFQ will specify 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 under the producer's direction, 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.

Direct performance incentives clearly vary with the choice of contractual form. Under a dayrate contract, which is essentially cost-plus, the producer will have a direct incentive to design an efficient drilling program but the rig will not be directly incented to exert a high level

⁵ Exceptions to single-well contracting tend to occur in large well-established fields where geologic uncertainty is low. For example, trade publications and interviews with industry participants have indicated that, in the large Barnett Shale gas field in East Texas, development wells are virtually guaranteed to find gas and producers there regularly sign long-term contracts with their rigs. See, for example, RigZone (2006).

⁶ An inefficiency may arise if firms are capable of enhancing relationship-specific learning through costly investments such as job training. In the absence of a long-term contract, firms may under-invest. Such an inefficiency may explain why some production companies and rigs use long-term contracts in fields such as the Barnett Shale where geologic uncertainty is low.

of effort. A footage contract, in contrast, places the full direct incentive on the rig. However, indirect performance incentives are important under both contract types. In the case of footage contracts, producers can use efficient well designs, backed by historically low drilling times, to obtain lower bids from drillers. As for dayrate contracts, interviews with industry participants have revealed that rig reputations are well-known by producers and that rigs 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. In an environment in which repeat contracting is common, this observability of effort can generate implicit performance incentives for the contractor (Corts 2007).⁷

2.3 Productivity and drilling time

This paper uses 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. 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 drill 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 in my empirical model: the prices charged by rigs rise during periods of high drilling activity, which will create spurious correlation between drilling cost and experience.

⁷ Corts and Singh (2004) assess the determinants of contract type in the offshore drilling industry, and the onshore data I use here support his conclusions. In section 6.5, I address the possibility that the learning analysis presented in this paper is confounded by changes in firms' choice of contractual form.

3. Firm-Specific and Relationship-Specific Learning

This paper considers learning that is both firm and relationship-specific. Firm-specific learning refers to improvements in a firm's productivity that are associated with increases in the firm's experience. This "standard" learning-by-doing effect has been widely documented in the empirical literature, beginning with Wright's (1936) and Alchian's (1963) studies of aircraft manufacturing. Relationship-specific learning, on the other hand, refers to productivity increases that depend not only on a firm's general experience but also its joint experience with the particular firms with which it works. These joint experience effects have received little attention, though McCabe (1996) finds evidence suggestive of relationship-specific learning in the construction of nuclear power plants: the productivity of primary construction contractors engaged in brief relationships with their utilities was lower than that of contractors in long-term relationships.⁸

In the drilling industry, mechanisms exist for learning along three dimensions: (1) producer-level firm-specific learning; (2) rig-level firm-specific learning; and (3) relationship-specific learning between producers and rigs working together. Producer-specific learning occurs because every well drilled into a field yields information regarding both the field's geology and which drilling procedures work well in that geology. For example, the optimal selection of drilling bits and drilling mud depends critically on the types of rock encountered. Producers' learning is therefore technical in nature and tends to be field-specific. This learning is well-recognized within the drilling industry, and several engineering case studies have documented how producers use past experience to reduce drilling times. See, for example, Brett and Millheim (1986) and Adeleye *et al.* (2004).

Because rigs are usually not involved in well design and planning, rig-specific learning is less technical in nature than is producer-level learning. Instead, rigs' learning comes from improved teamwork and developments in crew members' skills. For example, crews become more efficient at lowering drilling pipe into a hole, 90 feet at a time, after carrying out this same task on numerous wells in the past.

Finally, several mechanisms of relationship-specific learning are possible. The rig's crew may become familiar with the producer's particular drilling procedures, or the producer's 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

⁸ McCabe (1996) does not discern, however, whether the productivity differences in the data are driven by relationship-specific learning or by other utility-specific or relationship-specific heterogeneities.

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, particularly if they have dealt with drilling problems together in the past.⁹

The intuition behind relationship-specific learning has a parallel in recent theoretical work. Ellison and Holden (2008) develop a model in which a principal hires an agent to repeatedly take an action. The optimal action in each period is state-dependent, but the principal cannot communicate a complete contingent plan to the agent. Thus, in some states of the world, the agent may not take the optimal action. However, once a state has been realized and acted upon, the principal gains the ability to communicate the optimal action for that state, so that the agent can take that action when the state occurs again. In this way, the firms' performance improves as they accumulate experience working together.

4. Data

The central empirical challenge of this paper is to separate the impact of relationship-specific learning from that of firm-specific 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 oil and gas industry regulator. 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 67.7% of observations have a drilling time prior to this date, compared to 89.8% afterwards. I therefore focus my analysis on 1991-2005, during which there exist 106,946 TRRC observations with a recorded drilling time.¹⁰

⁹ These mechanisms suggest that relationship-specific learning occurs between the producer's company man and the rig's foreman and crew rather than the producer and the physical rig itself. The ideal empirical analysis would therefore use data on the duration of relationships between producer and rig personnel (the rig foreman in particular). However, I only possess data on relationships between producers and rigs, not personnel, so am measuring the true relationship of interest with error. This error may not be too severe given that rig foremen typically have multi-year spells with a single rig, but will nonetheless attenuate estimates of relationship-specific learning.

¹⁰ While the TRRC asks producers to report the date drilling began for all their wells, this reporting is not rigorously enforced. Beyond the missing data, 2.7% 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

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.¹¹ Unlike the TRRC data, the unit of observation in the 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.

The empirical analysis 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 four 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 6.1% of the TRRC dataset, have no corresponding record in the SB data and are therefore impossible to match. In section 6.5, I verify that the selective removal of these wells does not substantially impact the empirical analysis.

Second, 14.8% of the 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 ambiguous cases unmatched to avoid the risk of matching firms that are, in fact, distinct.

Third, 27.7% of the TRRC wells do not have a match because the SB data are not as comprehensive as the TRRC data: SB records 23.3% fewer drilling-weeks than does the TRRC. These match failures do not appear to be systematic; in particular, their incidence is not significantly correlated with wells' drilling times, the primary dependent variable of the analysis.¹² Finally, some non-unique matches occur when a producer employs multiple drilling

a single day. The incidence of these observations and those with missing drilling times is not correlated with the experience variables that I ultimately use in my analysis.

¹¹ Unfortunately, I do not observe the price charged or whether the rig is on a one-well or multi-well contract.

¹² Specifically, I regress a flag for whether each TRRC observation matched at least one SB observation on the log of the well's drilling time and a set of field X producer fixed effects. The point estimate on the log of drilling time is -0.0085—small in magnitude—with a standard error of 0.0060.

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, reducing the dataset by a further 20.4% of the original TRRC well count.

This matching process yields a dataset with 33,125 observations for which the producer and drilling rig are known. Of these wells, 7.7% are exploratory wildcats and are dropped because the field location is not recorded. In addition, because horizontal and directional wells are typically best-drilled with specialized rigs, I omit these wells, comprising 20.2% of the data, from my analysis. I also drop dry holes, comprising 14.3% of the remaining observations, because their drilling times can be artificially inflated if the producer keeps the rig on-site while it attempts to coax the well to flow.¹³

Finally, I drop all fields, producers, and rigs for which there is only one observation since tracking learning for such entities is not possible. The final matched dataset consists of 19,059 wells, spread over 1,354 fields, 704 producers, and 1,339 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 23.0 days, the average well depth is 9,036 feet, and 90% of the data lie between 4,650 feet and 13,800 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

¹³ While horizontal, directional, and dry holes are not used in the final dataset, I still "count" the fact that they were drilled when I calculate the experience variables for the associated field, producer, and rig. Although the field locations of wildcats are unknown, their drilling is included in the experience of the associated producer and rig.

contractor relationships into the analysis. In section 6, I examine how the results presented here are affected by taking relationship-specific learning into account.

5.1 Empirical framework

This paper follows the learning-by-doing literature by estimating a production function in which firms' past experience is the measure of human capital accumulation. The typical specification in the literature has the Cobb-Douglas form given by

$$y_{fpt} = (E_{fpt})^\alpha (K_{fpt})^\beta (L_{fpt})^\gamma v_{fpt} \quad (1)$$

where y_{fpt} measures the productivity (inverse drilling time) with which a well is drilled by producer p in field f at date t . E_{fpt} represents the producer's experience in field f at t ,¹⁴ K_{fpt} and L_{fpt} represent capital and labor inputs, and v_{fpt} represents factors unobserved to the econometrician; for example, geologic characteristics of field f . The magnitude of the coefficient α indicates the importance of learning-by-doing.

The capital and labor inputs to any given well are primarily determined by the well's drilling rig and its crew. While the analysis of section 6 will control for rig heterogeneity, the present analysis will instead assume that capital and labor inputs are constant within each producer and field. That is, I will ignore for now any changes in the identities of the rigs with which each producer contracts as well as any rig-level learning. The impact of these omissions on the estimated learning rates of field producers will be assessed in section 6.

Taking logs of (1) and rearranging to include field and producer fixed effects, the reference case specification for estimating producer-level learning is given by (2) below:

$$\log(DT_{fpt}) = f(E_{fpt}) + \gamma_f + \delta_p + \eta_t + \phi \mathbf{X}_{fpt} + \varepsilon_{fpt} \quad (2)$$

The dependent variable for each well is the logarithm of its drilling time DT_{fpt} , and the explanatory variable of primary interest is producer p 's experience in field f at time t , denoted by E_{fpt} . Field and producer-specific capital and labor inputs have been subsumed into field and producer fixed effects γ_f and δ_p . The field fixed effects also play an important role in controlling for the substantial geologic heterogeneity in drilling conditions across fields. Were producers to

¹⁴ This measure of experience will later be expanded to include measures of spillovers: experience by producer p in fields other than f and the experience of other producers in field f .

drill more frequently in fields that are “easier” to drill, the estimated learning effect would be biased downwards in the absence of these fixed effects.

Specification (2) also includes year fixed effects η_t to control for industry-wide technological change.¹⁵ I include variables X_{fpt} for well type (oil vs. gas vs. both) and well depth to control for within-field heterogeneity.¹⁶ X_{fpt} also includes month-of-year fixed effects to 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, including changes in oil and gas prices, discovery of new fields, and the identification of unexploited reserves in existing fields (through seismic imaging technology, for example). The identification assumption implicit in (2) is that producers do not drill wells more frequently in fields in which they have, for reasons unrelated to learning, specific skills or knowledge that make them more productive drillers. If this assumption is false, the estimate of (2) will be biased, yielding an inflated estimate of the learning effect. In practice, violations of this assumption seems unlikely to be substantial, given the numerous other factors that determine when and where a producer drills, such as its expectations about fields’ future oil and gas production. Nonetheless, I test the validity of this assumption by estimating a variant of (2) that includes fixed effects for field-producer interactions. While this specification will still be subject to biases arising from the omission of rig identities and rig-level learning, it will be robust to unobserved field-specific heterogeneity in producers’ drilling productivity.

5.2 Calculation of experience

I define E_{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 .¹⁷ 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

¹⁵ In alternative specifications, I use a polynomial function of time to control for technological change. Doing so does not substantially affect the estimated results.

¹⁶ Geologic heterogeneity is predominantly cross-field rather than within field. For example, regressing well depth on a full set of field fixed effects yields an R^2 of 0.88. Moreover, within-field depth variation is not correlated with producer-specific, rig-specific, or relationship-specific experience. For example, regressing the log of well depth on the log of each producer’s field-specific experience and a set of fixed effects for producer-field interactions yields an estimated coefficient on experience of -0.0024 with a standard error of 0.0020.

¹⁷ 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.

data. Were I to instead use this smaller dataset, I would understate each producer's experience, and variations in the retention of data across fields and producers would add noise to the calculation, causing attenuation bias in the estimation of (2).

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.¹⁸ I therefore have no means to calculate a cumulative experience measure. 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), Benkard (2000), and Thompson (2003) have demonstrated that experience effects decay with time as learning is "forgotten," supporting the importance of recent experience in determining productivity. In section 6.4, I discuss evidence of forgetting effects in the drilling industry.

As an alternative to measuring experience with E_{fpt} , the number of recently drilled wells, I have also calculated \hat{E}_{fpt} : the number of days of active drilling during the two years prior to t . It is not obvious which of these two variables is the more appropriate measure of experience. The intuition behind E_{fpt} is that learning by producers is technical and driven by the geologic information gained with each penetration rather than the accumulation of days of experience. However, if firms tend to learn more from mistakes than from successes, measuring experience in terms of time spent drilling, per \hat{E}_{fpt} , may be more appropriate. In the exposition of the estimation results, I will focus on those using E_{fpt} for two reasons: (1) this measure is used more frequently in the petroleum engineering literature; and (2) it ultimately yields more conservative estimates of the relationship-specific learning effect.

Measuring experience as the number of wells drilled within the past two years does, however, create potential for simultaneity bias that will cause an estimate of (2) to exaggerate the learning effect. Frequently, a rig will drill a series of wells for a producer 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 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 recently drilled. This simultaneity will cause a spurious negative correlation between drilling time and E_{fpt} in (2), exaggerating the estimated learning effect. I address this problem by instrumenting for E_{fpt} with \hat{E}_{fpt} . \hat{E}_{fpt} is not subject to the simultaneity problem: when wells are drilled back-to-back over two years or more, \hat{E}_{fpt} will

¹⁸ 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 6.5.

remain roughly constant as the drilling time per well decreases and the number of wells drilled increases.

To capture learning spillovers, both across fields and across producers, an alternative specification of (2) includes two new variables, E_{-fpt} and E_{f-pt} . These tally, respectively, the number of wells recently drilled by producer p in fields other than f and the number of wells recently drilled in field f by producers other than p . These variables are similar to those used by Thornton and Thompson (2001) in their study of learning spillovers in wartime shipbuilding, and summary statistics are presented in the upper section of table 3.¹⁹

5.3 Estimation results

To begin, I estimate (2) as written, without allowing for cross-field or cross-producer learning spillovers. Most studies of learning-by-doing model learning curves with a log-log (Cobb-Douglas) functional form, which in this setting implies that $f(E_{fpt}) = \beta \cdot \log(E_{fpt})$. Before taking this approach, I estimate (2) flexibly by fitting a cubic spline to $f(E_{fpt})$, instrumenting with a spline on \hat{E}_{fpt} . The results are plotted in figure 3. Drilling times are estimated to decrease by about 15% over the first 20 wells drilled by a field producer and then stay relatively constant over the remaining wells.

Figure 3 also plots the estimate of the log-log functional form. The point estimate of β is equal to -0.042 with a clustered standard error of 0.005.²⁰ This specification offers a reasonable fit to the spline, remaining within the 95% confidence interval of the spline estimate for the vast majority of observations. Most of the estimates reported in this paper are therefore of a log-log specification, though the results of estimating more flexible functional forms are also presented as robustness tests.

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. Deeper wells require more drilling time than shallow wells. Gas wells require more time to drill than oil wells, reflecting the heightened difficulty of managing wellbore pressures in the presence of gas. The year fixed effects indicate the presence

¹⁹ The logic behind instrumenting for E_{fpt} does not apply to the spillover variables because decreases in the drilling times of producer p in field f do not increase the number of wells that can be drilled in a two-year window by other producers in field f or in fields other than f .

²⁰ 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.

of some industry-wide technological improvement over the sample period, as drilling times decrease by approximately 20% 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 learning rate is larger than that of the reference case, though not substantially so. Column III replaces E_{fpt} with \hat{E}_{fpt} in (2), so that experience is measured as the number of days of drilling rather than the number of wells drilled. Here, the estimated learning effect is lower than that of the reference case: the point estimate of β is -0.020 rather than -0.042 and is estimated quite precisely with a standard error of 0.002. The source of this discrepancy may be that, if E_{fpt} is a better reflection of the process by which human capital is accumulated, then \hat{E}_{fpt} measures experience with error and the point estimate in column III reflects attenuation bias.

Column IV of table 5 addresses the possibility that producers drill wells more frequently in those fields that they are particularly good at drilling (for reasons other than learning) by including fixed effects for field-producer interactions. The estimated learning coefficient is -0.050, similar to the reference case estimate of -0.042. 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.

Regression V examines the importance of learning spillovers. Producers' experience in other fields appears to improve their productivity, though the magnitude of this effect is approximately one-half 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), who identify modest cross-firm 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.²¹

²¹ Conversations with industry participants have indicated that producers will sometimes include confidentiality clauses in their drilling contracts to prevent rig crews from sharing field-specific knowledge across producers.

6. Empirical Analysis: Rigs and Relationship-Specific Learning

The analysis presented thus far has followed traditional learning-by-doing studies in that it has omitted the impact of relationship-specificities on the learning process. The learning estimates presented in tables 4 and 5 attribute all learning effects to producers without allowing for 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 assess the degree to which the previous section's results misattributed learning effects solely to producers.

6.1 Empirical framework

I augment specification (2) 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 (3) below.

$$\log(DT_{fpri}) = f(\mathbf{E}_{fpri}) + \gamma_f + \delta_p + \eta_t + v_r + \boldsymbol{\varphi} \mathbf{X}_{fpri} + \varepsilon_{fpri} \quad (3)$$

In (3), 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 v_r , while γ_f , δ_p , and η_t denote field, producer, and year fixed effects as in (2). \mathbf{E}_{fpri} is now a vector of experience variables, and the reference case expands $f(\mathbf{E}_{fpri})$ per (4) below.²²

$$f(\mathbf{E}_{fpri}) = \beta_1 \cdot \log(E_{fpri}) + \beta_2 \cdot \log(E_{-fpri}) + \beta_3 \cdot \log(E_{f-pri}) + \beta_4 \cdot \log(E_{-pri}) + \beta_5 \cdot \log(E_{pri}) \quad (4)$$

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, each of these three variables is taken as the number of wells drilled within the past two years, and E_{fpri} is instrumented with \hat{E}_{fpri} , the number of drilling days accumulated by producer p in field f within the past two years.

The fourth term in (4) 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

²² The estimation of alternative functional forms for $f(\mathbf{E}_{fpri})$, for example a CES function, is discussed in section 6.5.

data as this avoids understating each rig’s experience. I define E_{prt} to be the number of weeks rig r was actively drilling for p within the two years prior to t and define E_{-prt} similarly. Each rig’s experience is measured in units of time rather than wells because rig-level learning occurs through the repetition of tasks and the accumulation of interactions with the producer, both of which should be functions of time. Moreover, this calculation is necessitated by the data, since Smith Bits tracks drilling activity in rig-weeks rather than well-by-well. Summary statistics for E_{-prt} and E_{prt} are indicated in the lower section of table 3.

Relationship-specific learning is important to the extent that β_5 is more negative than β_4 . These coefficients are identified through two sources of variation: (1) changes in the producer to which a given rig is contracted; and (2) the employment of multiple rigs (either simultaneously or in series) by a producer. If rigs have producer-specific characteristics (independently of human capital acquired through learning) and are likely to have longer relationships with producers to which they match well than those to which they do not, then this variation will not be exogenous and the estimate of β_5 , the coefficient on relationship-specific learning, will be biased downwards. Two institutional features of the drilling industry, however, suggest that match specificities are unlikely to be a serious concern. First, as noted in section 2.2, the rig equipment itself is generally not field or producer-specific, apart from the rig’s depth rating. Second, most producer-rig contract terminations in the data appear to be driven by a lack of additional drilling by the producer, with any rig, rather than by poor performance. When a producer releases a rig, it is rare that the rig is replaced with a new rig; instead, the producer simply ceases drilling, indicating that the relationship ended because the producer had no additional work to offer the rig. Specifically, only 12.8% of terminations are followed by the hiring of another rig by the producer within four weeks. These facts suggest that performance is unlikely to drive tenure, at least on average, in this setting. Nonetheless, I examine whether producer-rig specificities drive the estimate of (3) by adding fixed effects for producer-rig pairs to the specification. With the inclusion of these fixed effects, identification of relationship-specific learning comes only from variations in joint experience within each rig-producer pair.²³

²³ 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.

6.2 Primary estimation results

Column I of table 6 reports the results of estimating relationship-specific learning per equations (3) and (4). The estimated coefficient on $\log(E_{prt})$ —joint experience between a rig and a producer—is -0.025, 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 10%. However, the estimated coefficient on $\log(E_{-prt})$ is only -0.010, implying that a rig that frequently changes producers during a year can expect to decrease its drilling times by only 3.9%. Thus, on average, rigs with stable contracting relationships improve their productivity more than twice as quickly as rigs that frequently change contracting partners. Moreover, the difference between the coefficients on $\log(E_{prt})$ and $\log(E_{-prt})$ driving this result is statistically significant: an F-test rejects pooling with a p-value of 0.003.

In addition, the point estimate corresponding to learning by field producers is only -0.030 in this specification. This point estimate is lower in magnitude than was reported in column V of table 5, when the impact of producers' relationships with rigs was not considered in the regression. Thus, investigating learning using only the experience of producers overestimates the contribution of their stand-alone experience to observed productivity improvements.

Column II of table 6 adds fixed effects for producer-rig pairs to the specification and also removes from the data all pairs for which there is one observation. The estimate of relationship-specific learning is not significantly affected by the addition of these fixed effects: the new point estimate on $\log(E_{prt})$ is -0.023 rather than -0.025. The difference between the coefficients on $\log(E_{prt})$ and $\log(E_{-prt})$ is still statistically significant with a p-value of 0.026.²⁴ This result is consistent with a limited effect of producer-rig match specificities on relationship durations.

I use the estimated coefficients to obtain an estimate of the cost savings obtained through relationship-specific learning. In a counterfactual in which joint experience yields the same productivity benefit as stand-alone experience (that is, β_5 equals β_4 in (4)), the average drilling time in my sample would be increased by 3.8%, equal to 0.88 days at the sample average drilling time of 23.0 days. This efficiency gain is of comparable magnitude to that obtained on average from producers' stand-alone learning: 6.4% (1.5 days).²⁵

²⁴ The coefficient on $\log(E_{-prt})$ is still identified in the presence of producer-rig fixed effects because rigs sometimes have multiple employment “spells” with a single producer, and E_{-prt} will be different in each spell.

²⁵ The 3.8% figure is equal to the sample average of $\exp((\beta_4 - \beta_5) \cdot \log(E_{prt})) - 1$, and 6.4% is the sample average of $\exp(-\beta_1 \cdot \log(E_{prt})) - 1$.

At the 2005 rig dayrate of approximately \$11,000 per day for a well of average depth, the 0.88 day efficiency gain attributed to relationship-specific learning translates to an average reduction in rig rental cost of \$9,700 per well. Is this magnitude sufficient to drive firms to try to maintain long-term relationships? The data indicate that it is rare for a producer to switch rigs in the middle of a drilling program—the vast majority of contract terminations appear to be caused by a lack of new drilling work rather than a desire to change contractors. Section 7 investigates this pattern in the data more carefully in an effort to distinguish whether this contracting pattern is driven by relationship-specific learning or by other explanations such as switching costs or *ex ante* match specificities.

6.3 Mechanisms behind relationship-specific learning

This section examines more closely the mechanisms behind the observed relationship-specific learning. Is this learning driven by rigs' increasing familiarity with the characteristics of the fields that are drilled by their producers, familiarity with the particular procedures and personnel used by their producers, or both? I attempt to disentangle these mechanisms by decomposing each rig's experience into the following field-specific and non-field-specific components:²⁶

- (1) E_{-f-pri} : experience with producers other than p in fields other than f
- (2) E_{f-pri} : experience with producers other than p in field f
- (3) E_{-fpri} : experience with producer p in fields other than f
- (4) E_{fpri} : experience with producer p in field f

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 and report the results in the third column of table 6. The estimated coefficients suggest that both field-specific and producer-specific experience are important determinants of rigs' drilling productivity. The estimated coefficient on $\log(E_{fpri})$ is more than twice that on any of E_{-f-pri} , E_{f-pri} , or E_{-fpri} and is statistically distinct from each at the 5% level. The point estimates

²⁶ 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 a rig is active. I therefore estimate each rig's field-specific experience using a two-step procedure. 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. I then multiply this fraction by the total number of weeks the rig has been active during the past two years, taken from the SB data.

imply that a rig that works for the same producer in the same field for a year can expect a 10% increase in drilling productivity. However, were the rig to then switch producers (or fields), its productivity would on average be only 3% (or 4%) larger than that of a rig with no experience at all.

In column IV of table 6, I repeat the specification of column III but include fixed effects for the triple interaction of field, producer, and rig identifiers. Similar to the producer-rig fixed effects applied in columns I and II, these terms address the possibility that the column III results are driven by match-specificities between fields, producers, and rigs. Including these fixed effects removes from the sample nearly 5000 observations that are associated with a field-producer-rig for which I observe only one well and requires 3,299 fixed effects. This reduction in degrees of freedom hinders inference, as evidenced by the increase in the standard errors of the estimated coefficients. While the magnitudes of the coefficients in column IV fall roughly in line with those in column III, the coefficient on $\log(E_{fprt})$ is no longer statistically distinguishable from those of E_{f-prt} , E_{f-prt} , or E_{fprt} .

Thus, while conclusions drawn from this set of results must be tempered, they do provide some suggestive evidence that knowledge acquired through work in a field with one producer is not fully portable to other producers in the same field. This result is consistent with the views expressed to me by industry participants that the building of a personal relationship between producer and rig personnel is important, and with the evidence presented in section 5 indicating that cross-producer learning spillovers are, at best, small. In addition, the lessons learned by a rig in one field do not appear to be completely transferable to a different field, even if the rig continues to work for the same producer. This result seems likely to be driven by variations in producers' drilling procedures or personnel across fields.

6.4 Forgetting effects

This section investigates whether experience effects in the drilling industry decay over time, consistent with institutional “forgetting” of knowledge. Specifically, 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), Benkard (2000), and Thompson (2003) that identified forgetting effects in aircraft manufacturing and shipbuilding. I first examine producers' forgetting in the setting of section 5, in which rig and relationship-specific learning are not considered, paralleling this literature. I then evaluate producers'

forgetting within the full model of sections 6.1 and 6.2 to investigate the extent to which forgetting effects can be explained by relationship-specific learning.

Thus far, I have defined E_{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 define E_{fpt} as a function of a decay parameter δ , per expression (5) below, in which $N_{fp\tau}$ denotes the number of wells drilled by producer p in field f on date τ .

$$E_{fpt}(\delta) = \sum_{\tau=t-730}^t e^{\delta(t-\tau)/365} \cdot N_{fp\tau} \quad (5)$$

For negative values of δ , wells drilled on dates long before t carry less weight in $E_{fpt}(\delta)$ than do wells drilled near date t . I estimate δ by inserting (5) into learning specification (2), which does not allow for rig or relationship-specific learning, yielding expression (6) below.

$$\log(DT_{fpt}) = \beta \cdot \log(E_{fpt}(\delta)) + \gamma_f + \delta_p + \eta_t + \boldsymbol{\varphi} \mathbf{X}_{fpt} + \varepsilon_{fpt} \quad (6)$$

I estimate (6) using nonlinear least squares. As in section 5, I instrument for experience using the number of days producer p actively drilled in field f during the two years prior to date t .²⁷ I obtain a point estimate of δ equal to -0.813 with a clustered standard error of 0.420, consistent with the presence of forgetting. While this result suggests, at first glance, that $\hat{\delta}$ is statistically significant (with a p-value of 0.053), the proper statistical test for forgetting must take into account the negative correlation between $\hat{\delta}$ and $\hat{\beta}$. The point estimate of β is -0.051, larger in magnitude than in the reference case results without relationship-specific learning (table 5, column I). This increase in magnitude occurs because learning is now a function of depreciated experience rather than total experience. Because the estimates of δ and β are linked, I test for forgetting by testing whether these estimates are jointly different from those reported in the reference case, in which $\delta = 0$ and $\beta = -0.042$. I find that the null hypothesis of no forgetting can be rejected with only a p-value of 0.149, rather than at nearly the 5% level.

The estimated rate of experience depreciation is somewhat large: the point estimate of δ implies that a well drilled one year ago makes a contribution to experience that is only 44% 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

²⁷ The estimation also instruments for the derivative of experience with respect to δ using the derivative of drilling time-based experience with respect to δ .

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. I investigate this possibility by augmenting (6) with rig fixed effects and variables measuring rig and relationship-specific experience, per equations (3) and (4). While the new point estimate of δ from this specification is still negative and quite large in magnitude at -1.43, it is estimated quite imprecisely: testing against a null hypothesis of no forgetting yields a p-value of only 0.282.²⁸ This result suggests that losses of relationship-specific capital between lead firms and contractors may be one of the mechanisms behind forgetting effects documented by other studies.

6.5 Alternative specifications and robustness

Functional form: Table 7 presents the results of estimating functional forms that are alternatives to the Cobb-Douglas specification given by equations (3) and (4). One such form that nests Cobb-Douglas as a special case is the constant elasticity of substitution (CES) relation (7):

$$f(E_{fppt}) = \beta_1 \cdot \log(E_{fppt}) + \beta_2 \cdot \log(E_{-fppt}) + \beta_3 \cdot \log(E_{f-ppt}) + (\delta/\rho) \cdot \log((1-\alpha) E_{-ppt}^\rho + \alpha E_{ppt}^\rho) \quad (7)$$

In (7), the importance of relationship-specific learning is evaluated by testing whether α is greater than $\frac{1}{2}$. The Cobb-Douglas case that has been considered thus far is equivalent to the limit of (7) as ρ approaches zero. As an alternative to Cobb-Douglas, I first consider in column II of table 7 another special case of (7) in which ρ is fixed at unity: perfect substitution. Using nonlinear least squares, I find that the point estimate of α under perfect substitution is 0.928 and statistically distinct from 0.5 with a standard error of 0.082, implying a strong relationship-specific learning effect. Together with the point estimate of δ of -0.033, this result implies that a rig that works for a single producer over the course of a year will improve its productivity by 13% on average, relative to a 5.2% improvement for a rig that frequently changes producers. That is, after a year, relationship-specific learning drives a productivity improvement 2.5 times as great as that driven by general learning, nearly the same ratio that was estimated in the Cobb-Douglas specification.

Column III of table 7 reports the results of estimating (7) without any restrictions on ρ . The point estimate of ρ is 0.263 with a standard error of 0.432, clearly failing to reject Cobb-Douglas but rejecting perfect substitution at the 10% level. The point estimate of α is 0.712 and

²⁸ When allowing for forgetting, the importance of rig's producer-specific experience is still statistically distinct from that of rig's stand-alone experience at the 1% level.

statistically distinct from 0.5 at the 10% level, again supportive of relationship-specific learning effects.

In column IV, I estimate a flexible functional form in which E_{fpt} , E_{-fpt} , and E_{f-pt} enter the specification as cubic splines, while leaving the variables tied to rig experience, E_{-prt} and E_{prt} , in a parameterized Cobb-Douglas form. That is, I model:

$$f(\mathbf{Experience}_{fprt}) = s_1(E_{fpt}) + s_2(E_{-fpt}) + s_3(E_{f-pt}) + \beta_4 \cdot \log(E_{-prt}) + \beta_5 \cdot \log(E_{prt}) \quad (8)$$

The estimates of β_4 and β_5 from (8), presented in column IV, are similar to those of the reference case, indicating that the relationship-specific learning result is not driven by the parameterization in (4). Going a step further, I also estimate a version of (8) in which all five forms of experience enter through splines. The estimated functions $s_4(E_{-prt})$ and $s_5(E_{prt})$ are plotted in figure 4, which shows that relationship-specific experience has a stronger effect on performance than does general experience across the full domain of data. Moreover, the two curves are point-wise statistically distinct at the 5% level for experience values between 1 and 86 weeks, beyond which point the sample size becomes too small to separate them.

Finally, in column V of table 7, I replace the fourth term in (4), $\beta_4 \cdot \log(E_{-prt})$, with $\beta_4 \cdot \log(E_{rt})$: the total experience of rig r in the two years prior to date t . In this case, relationship-specific learning is no longer measured by the difference between β_4 and β_5 but by the difference between β_5 and zero. I find that the point estimates of both β_4 and β_5 are equal to -0.018 and that both are significantly different from zero at the 1% level. This result accords well with the reference case specification, as it implies that a rig that maintains a stable relationship will improve its performance twice as quickly as a rig that frequently changes producers.

Contract type: Could 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? The analysis to this point has not taken the firms' choice of dayrate or footage contract into account. If firms tend to switch from dayrate to footage contracts (or vice-versa) over the course of a relationship, the performance improvements that I observe may be driven by changes in incentives rather than learning.

I address the potential impact of changes in contract type by taking advantage of the fact that most producer-rig pairs do not change contractual form during the sample period. 76.4% of all observations are associated with a producer-rig pair that either always uses a dayrate contract or always uses a footage contract. 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 II of table 8. I still find a strong and statistically significant effect of joint experience on drilling time. The coefficient on $\log(E_{prt})$ is -0.022, very similar to that found when the same regression was run on the full sample (column II of table 6). The joint experience effects evident in the data do not appear to be driven by changes in contract type.²⁹

Measurement and data issues: 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 population, it is possible that this selection on the dependent variable may bias the results. I address this concern by estimating (3) and (4) with data only for wells that are at least 8000 feet deep (12,581 observations). Such wells are essentially impossible to drill in less than one week, and estimation with this sub-sample neutralizes the potential selection problem. Results, presented in column III of table 8, are very similar to those obtained from the full sample, shown in column I. The difference between the coefficients on $\log(E_{-prt})$ and $\log(E_{prt})$ is statistically significant with a p-value of 0.022.

The results reported thus far measure E_{fpt} , E_{-fpt} , and E_{f-pt} as the number of wells drilled within two years of date t . Table 8, column IV presents the results of estimating an alternative specification in which each of these variables is measured as the number of days of drilling within two years of date t . In this specification, the importance of the producer's experience, both within and outside of its current field, is diminished relative to the reference case model in column I. This decrease in the learning estimates parallels that discussed in section 5.3. Here, the alternative specification also slightly strengthens the estimated relationship-specific learning effect: the estimated coefficients on $\log(E_{-prt})$ and $\log(E_{prt})$ are -0.009 and -0.027, respectively, as compared to -0.010 and -0.025 in the reference case. Column V indicates that these results are robust to the inclusion of fixed effects for producer-rig interactions to the specification.

²⁹ I have run a related regression in which I parse the effect of joint experience into that for producer-rig pairs using dayrate contracts and that for pairs using 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 contracting—these wells tend to be more geologically challenging than those drilled under footage contracts and likely present greater scope for learning.

Finally, I verify that the results are robust to changes in the length of time over which I calculate the experience variables. The results in columns II, III, and IV of table 9 calculate experience over one, three, and five years, respectively, rather than the reference case of two years. In all three formulations, the coefficients on rigs' general and producer-specific experience change little and remain statistically distinct at at least the 5% level. These estimates remain precise despite the fact that, at longer windows of experience, the early years of the data must be dropped because the Smith Bits data only exist back to 1989.

7. Empirical Analysis: Relationship Persistence

In this section, I empirically examine 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 use rigs with which they have substantial prior experience? If so, is this preference driven by learning or by other factors?

It is clear from the data that firms do generally maintain their relationships. Producers rarely change rigs in the middle of a drilling program: relationships tend to end only when a producer runs out of work. This fact alone, however, can be supported by explanations other than learning and in particular by the presence of switching costs. I therefore test for relationship persistence by focusing on instances in which a producer has two rigs drilling for it in the same county. When the producer releases one of these two rigs, I ask whether the rig that is released is that with the least producer-specific experience. This last in-first out (LIFO) pattern would be consistent with firms' maximization of the benefits of relationship-specific learning and would also be difficult to explain using switching costs alone since the test is conditioned on a switching cost being paid.³⁰ This pattern may, however, be consistent with the presence of other *ex ante* match specificities as discussed below.

I execute this analysis using the original SB dataset, prior to its match with the TRRC data. There are 323,730 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. Week-to-week, rigs maintain their relationship with their producer 89.5% of the time. Rigs change producers in 7.4% 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.

³⁰ A LIFO pattern may result if switching costs are heterogeneous, a possibility I examine below.

I define all instances in which a producer has two rigs drilling for it in the same county as a “pair” and use these pairs as the unit of observation in my analysis.³¹ There are 907 unique pairs in the data, and with two rigs per pair, there exist 1,814 total observations, spread over 821 unique rigs and 531 unique county-producer combinations.

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 relationship-specific learning analysis of section 6. I then test for a systematic relationship between each rig’s experience and the identity of the rig that is released first: as the difference in producer-specific experience between the two rigs grows, does it become more likely that the less-experienced rig is released first?

Figure 5 illustrates the evidence of relationship persistence in this sample. The horizontal axis plots the absolute value of the difference (in logs) of producer-specific experience between the two rigs in each pair. Thus, points plotted on the right side of the graph represent observations in which the two rigs have very different levels of producer-specific experience. All observations are grouped into bins of width 0.2, and the vertical bars indicate the number of pairs in each bin. Each data point indicates the percentage of pairs within each bin for which the less-experienced rig was the first to exit. There exists a clear systematic pattern in the data: as the difference in specific experience between the two rigs in each pair grows larger, it becomes more and more likely that the less-experienced rig will exit first. This pattern is consistent with firms’ recognition of relationship-specific learning’s benefits.

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 (9) below, in which $Experience_{i1}$ denotes the producer-specific experience of rig 1 in pair i .

$$\Pr(ExitFirst_{Pair\ i, Rig\ 1}) = \frac{\exp(\beta \cdot \log(Experience_{i1}))}{\exp(\beta \cdot \log(Experience_{i1})) + \exp(\beta \cdot \log(Experience_{i2}))} \quad (9)$$

The results of this regression are reported in column I of table 10: rigs with more producer-specific experience are significantly less likely to exit first. The estimated marginal effect of -0.061 implies that, in a pair consisting of a rig with 12 months of experience and a rig

³¹ 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 poor data tracking by Smith Bits.

with 1 month of experience, the less experienced rig has a 63.7% probability of being the first to exit.

Column II of table 10 presents the results of estimating (9) 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 reflects the fact that the general experience of a rig does not provide productivity benefits that are producer-specific. While a highly experienced rig may be more 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.

Of course, there exist alternative explanations behind the LIFO pattern exhibited in figure 5 and column I of table 10. First, it may be that some rigs have lower switching costs than others and therefore change jobs frequently. A rig market in which some rigs are "switchers" and others are "stayers" would generate a LIFO pattern even without a relationship-specific learning effect. Second, *ex ante* match specificities would cause producers to first hire those rigs to which they match best and to release those rigs last.

To rule out heterogeneity in rig switching costs, I re-estimate (9) while including a set of rig fixed effects. I do so using a linear probability model, since including rig fixed effects in a conditional logit is likely to lead to an incidental parameters problem that will cause the estimate of β to be inconsistent (Neyman and Scott 1948, Lancaster 2000). Fortunately, the baseline results do not appear to be sensitive to model choice: as indicated in column III of table 10, estimating a linear probability model with the log of producer-specific experience and group fixed effects as covariates yields a marginal effect very close to that of the conditional logit. When rig fixed effects are added to the specification, I still find a strong systematic LIFO effect. As shown in column IV, the marginal effect is -0.067, compared to -0.059 in column III, and still statistically significant at the 5% level. Rig heterogeneity is not driving the LIFO result.

Taking a step further and adding fixed effects for producer-rig interactions would eliminate the influence of match specificities on this result. Unfortunately, within the 1814 observations in the sample there are 1,488 unique producer-rig combinations. The limited sample variation remaining after including these fixed effects precludes inference, as indicated in column V of table 10. The standard error of the estimated marginal effect nearly triples, relative to column IV, and this regression provides no evidence either for or against the LIFO pattern.

Supplemental evidence, however, supports the hypothesis that the relationship persistence is not driven by matching. First, I use information on rig depth ratings and well depths to assess how the matching of rigs' attributes to producers' wells affects producers'

hiring. For each rig in each pair, I calculate the absolute difference between the rig's depth rating and the average depth of the wells it drills.³² Across pairs, I find that the distribution of this "depth difference" for the first rig to enter each pair is very similar to that of the rig that enters last: a Kolmogorov-Smirnov test for the equality of these distributions fails to reject equality with a p-value of 0.582. This result indicates that the two rigs in each pair are equally well-matched in depth rating to their wells.

Further, I investigate pre-existing performance differences between the rigs in each pair by comparing the drilling times of the first well drilled by each rig within that pair. As with the depth differences, the data fail to reject equality of the initial performance of the rig that enters first with that of the rig that enters second: the Kolmogorov-Smirnov p-value is 0.886.³³ This similarity between both the initial performance and the depth rating of the two rigs in each pair suggests that the LIFO pattern in the data is not driven by pre-existing match specificities between producers and rigs.

The evidence is therefore consistent with a recognition by producers that maintaining long-term relationships helps to maximize the productivity benefits of relationship-specific learning. I cannot, however, rule out the possibility that these results are purely driven by friendships formed between the personnel of each firm. Personal relationships are important in the Texas drilling industry, and friendships may develop over the course of a long business relationship that neither party would want to break. Of course, friendship is likely to play a role in driving relationship-specific learning in the first place, by facilitating communication between producers and rigs. The value of empirically differentiating the learning and friendship stories is therefore not obvious, even if the necessary data were available.

8. Conclusions

This paper demonstrates that relationship-specific learning can be an important driver of productivity improvement and play a role in firms' contracting decisions. I find that a drilling rig that accumulates experience with one producer improves its productivity more than twice as quickly as a rig that frequently changes contracting partners. As a consequence, producers and rigs have a strong incentive to maintain their relationships, and the data demonstrate that

³² The average depth difference is 2,898 feet, with a standard deviation of 2,675 feet.

³³ On average, the first-well drilling time of the first rig to enter is actually 4.8% higher than that of the second rig (this difference is not statistically significant). This fact likely reflects producer-specific learning: by the time the second rig enters, the producer has learned from the wells it drilled with the first rig and improved its drilling times accordingly.

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 investments plausibly increase the rate of learning, the accumulated knowledge that results is also a form of relationship-specific capital. Thus, when learning rates may be enhanced through costly actions, it may be in firms' interests to develop contracting arrangements that alleviate *ex post* bargaining problems and promote efficient investment. Relationship-specific learning may therefore play a role in promoting long-term contracts and vertical integration.

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 drilling may be related to the competitive nature of production from a common pool resource. Further study with richer data may be fruitful in illuminating the drivers of learning spillovers and in assessing the importance of their macroeconomic effects.

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.
- Arrow, Kenneth J., "The Economic Implications of Learning by Doing," *Review of Economic Studies* 29 (Jun., 1962), 155-173.
- Becker, Gary S., *Human Capital; a Theoretical and Empirical Analysis, with Special Reference to Education*, New York: Columbia University Press (1964).
- Benkard, C. Lanier, "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review* 90 (Sep., 2000), 1034-1054.
- Benkard, C. Lanier, "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft," *Review of Economic Studies* 71 (2004), 581-611.
- Besanko, David, Ulrich Doraszelski, Yaroslav Kryukov, and Mark Satterthwaite, "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics," Harvard Institute of Economic Research Discussion Paper Number 2128 (Feb., 2007).
- 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).
- Cabral, Luis M.B. and Michael H. Riordan, "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica* 62 (Sep., 1994), 1115-1140.
- 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.

- Ellison, Glenn and Richard Holden, "A Theory of Rule Development," MIT working paper (Feb., 2008).
- Gibbons, Robert and Michael Waldman, "A Theory of Wage and Promotion Dynamics Inside Firms," *Quarterly Journal of Economics* 114 (Nov., 1999), 1321-1358.
- Gibbons, Robert and Michael Waldman, "Task-Specific Human Capital," *American Economic Review Papers and Proceedings* 94 (May, 2004), 203-207.
- 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.
- 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.
- Lancaster, Tony, "The Incidental Parameter Problem Since 1948," *Journal of Econometrics* 95 (2000), 391-413.
- Lucas, Robert E., "On the Mechanics of Economic Development," *Journal of Monetary Economics* 22 (1988), 3-42.
- Managing Automation*, "Boeing's Big Supply Chain Wager" (17 May, 2007).
- McCabe, Mark J., "Principals, Agents, and the Learning Curve: The Case of Steam-Electric Power Plant Design and Construction," *Journal of Industrial Economics* 44 (Dec., 1996), 357-375.
- 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.

- Poletaev, Maxim and Chris Robinson, "Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984-2000," *Journal of Labor Economics* 26 (2008), 387-420.
- Prendergast, Canice, "The Role of Promotion in Inducing Specific Human Capital Acquisition," *Quarterly Journal of Economics* 108 (May, 1993), 523-534.
- RigZone, "Dune Energy Closes Acquisition of Barnett Shale Properties for \$7.0 million" (16 Oct., 2006). http://www.rigzone.com/news/article_pf.asp?a_id=37135
- Spence, Michael A., "The Learning Curve and Competition," *Bell Journal of Economics* 12 (Spring 1981), 49-70.
- Stokey, Nancy L., "Learning by Doing and the Introduction of New Goods," *Journal of Political Economy* 96 (Aug., 1988), 701-717.
- Thompson, Peter, "How Much Did the Liberty Shipbuilders Forget?" Working paper, Florida International University, Miami (2003).
- 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.
- 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., *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

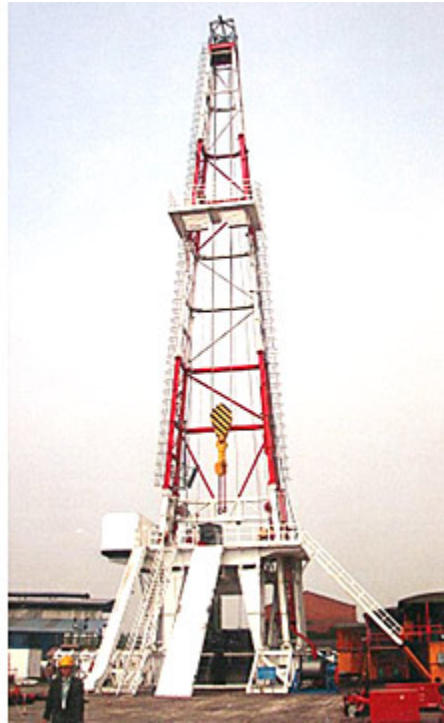
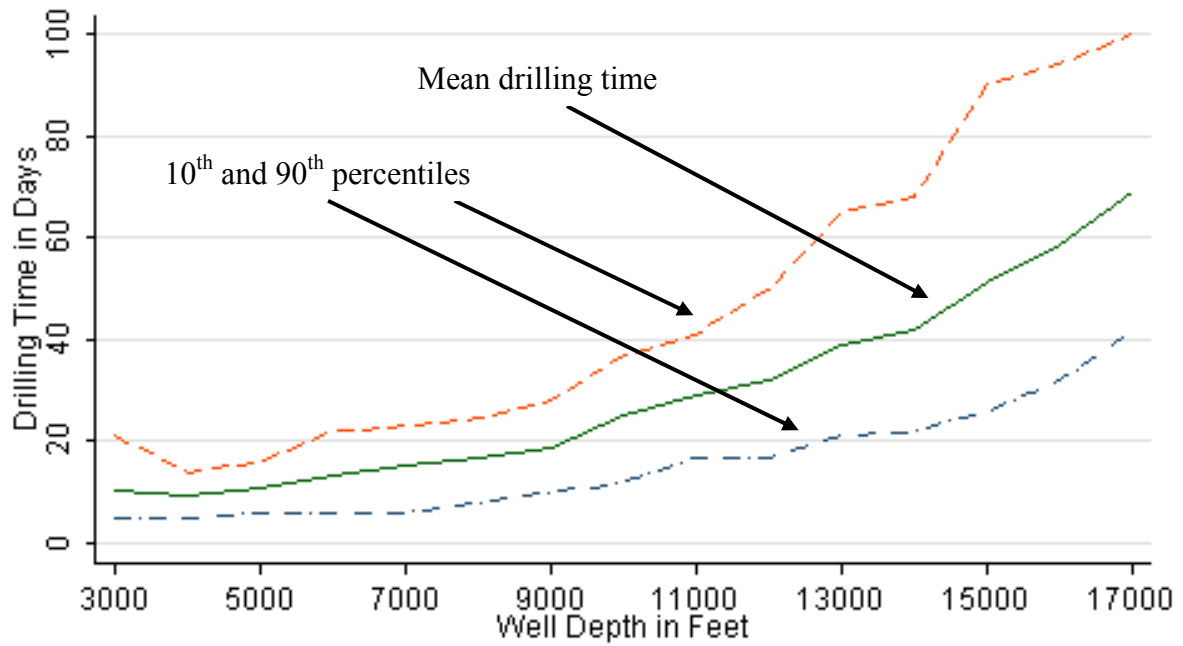
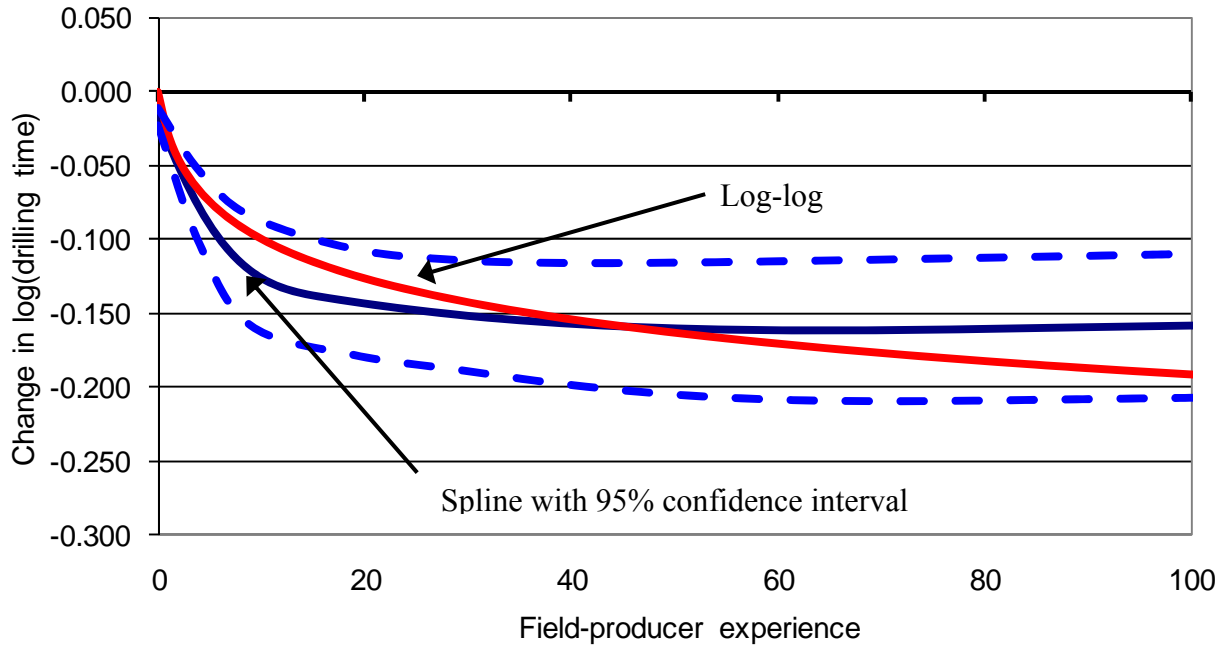


Figure 2: Drilling times vs. well depths



Note: Observations grouped into depth “bins” of 1000 feet
Horizontal axis excludes highest and lowest 1% of depths

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



Note: Plot includes 96% of data

Figure 4: Estimated splines for rigs' general and relationship-specific learning

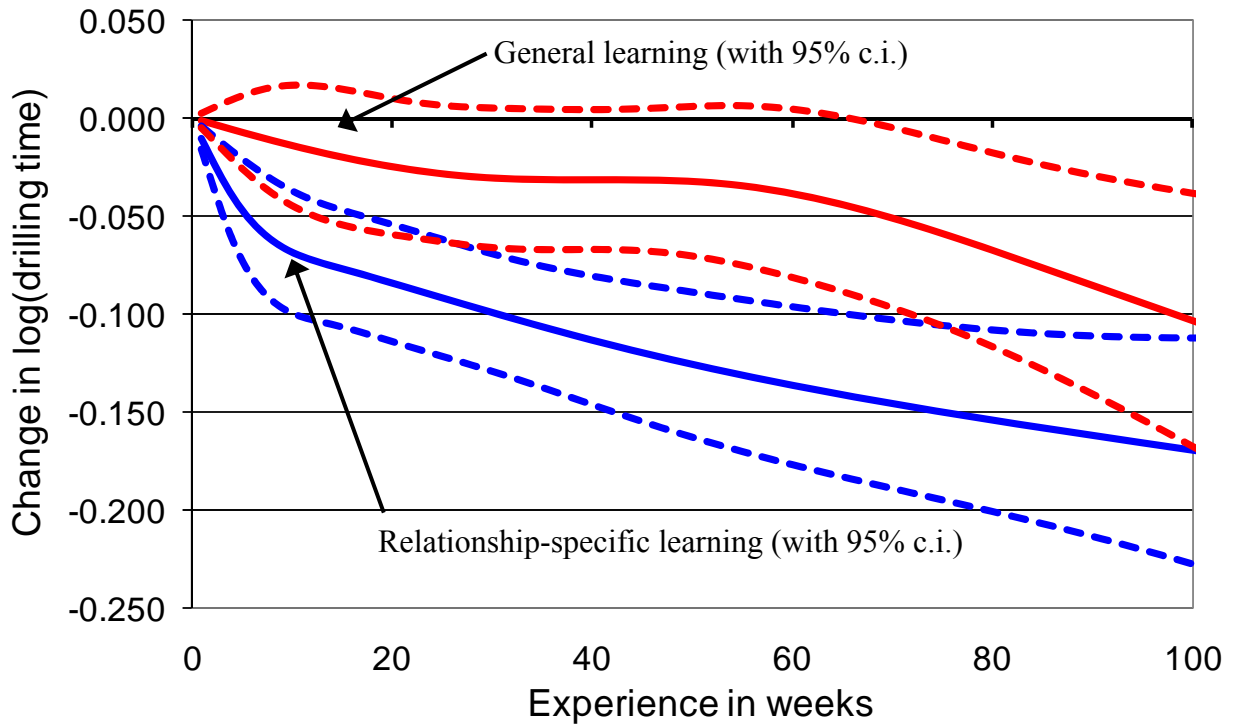


Figure 5: Likelihood that the least experienced rig is the first to change producers vs. the within-pair difference in rigs' producer-specific experience




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	10	14.1	784
Number of wells per producer	2	3	7	20.5	27.1	630
Number of wells per rig	2	4	8	19	14.1	157
Number of producers per field	1	1	2	3	2.9	54
Number of fields per producer	1	1	3	6	5.5	124
Number of rigs per driller	1	1	3	6	7.7	194

Table 2: Sample summary statistics

	Number of observations	Min	Median	Mean	Std. Dev.	Max
Drilling time (days)	19059	2	18	23.0	19.2	179
Well depth (feet)	19059	631	9000	9036	2817	23000
Gas well (0/1 dummy)	19059	0	1	0.620	0.485	1
Oil and gas well (0/1 dummy)	19059	0	0	0.001	0.036	1

Table 3: Summary statistics of experience variables

	Number of observations	Min	Median	Mean	Std. Dev.	Max
Number of wells drilled during the past two years in:						
Same field, same producer	19059	1	7	22.0	51.0	711
Different field, same producer	19059	1	46	125.7	169.6	1098
Same field, different producer	19059	1	10	69.6	171.6	1813
Number of weeks of drilling within past two years by:						
Same rig, different producer	19059	1	29	34.1	29	105
Same rig, same producer	19059	1	14	27.7	30.6	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 4: Regression results for learning by field producers
Dependent variable is log(drilling time)

Variable	Point estimate	Standard error	Variable	Point estimate	Standard error
log(field-producer exper.)	-0.042	(0.005) ***	1992	-0.040	(0.026)
log(well depth)	1.179	(0.075) ***	1993	-0.016	(0.022)
Gas well	0.058	(0.026) **	1994	-0.062	(0.030) **
Oil and gas well	0.239	(0.081) ***	1995	-0.060	(0.035) *
February	-0.021	(0.014)	1996	-0.038	(0.032)
March	-0.030	(0.013) **	1997	-0.026	(0.036)
April	-0.048	(0.014) ***	1998	-0.035	(0.039)
May	-0.045	(0.013) ***	1999	-0.089	(0.039) **
June	-0.062	(0.014) ***	2000	-0.072	(0.037) *
July	-0.046	(0.014) **	2001	-0.022	(0.039)
August	-0.034	(0.014) **	2002	-0.087	(0.036) **
September	-0.016	(0.014)	2003	-0.116	(0.040) ***
October	-0.053	(0.014) ***	2004	-0.191	(0.043) ***
November	-0.030	(0.014) **	2005	-0.204	(0.042) ***
December	-0.033	(0.015) **			

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)

	I	II	III	IV	V
	Reference case (Table 4)	Experience not instrumented	Drilling time experience	Field-producer fixed effects	Learning spillovers
Log of experience with:					
Same field, same producer	-0.042 *** (0.005)	-0.053 *** (0.005)	-0.020 *** (0.002)	-0.050 *** (0.007)	-0.041 *** (0.005)
Different field, same producer	-	-	-	-	-0.022 *** (0.008)
Same field, different producer	-	-	-	-	-0.001 (0.007)
Field X producer dummies	N	N	N	Y	N
Number of observations	19059	19059	19059	19059	19059

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, month and year fixed effects, and field and producer fixed effects

Table 6: Regression results for relationship-specific learning
Dependent variable is log(drilling time)

	I	II	III	IV
Log of experience with:	Relationship-specific learning	Producer-rig fixed effects	Rig-producer-field specificities	Rig-producer-field fixed effects
Same field, same producer (<i>Experience_{fpt}</i>)	-0.030*** (0.006)	-0.027*** (0.008)	-0.027*** (0.007)	-0.032* (0.018)
Different field, same producer (<i>Experience_{-fpt}</i>)	-0.022*** (0.008)	-0.003 (0.012)	-0.022*** (0.008)	-0.015 (0.018)
Same field, different producer (<i>Experience_{f-pt}</i>)	-0.005 (0.005)	0.003 (0.008)	-0.005 (0.005)	0.005 (0.014)
Same rig, different producer (<i>Experience_{-prt}</i>)	-0.010** (0.004)	-0.005 (0.006)	- -	- -
Same rig, same producer (<i>Experience_{prt}</i>)	-0.025*** (0.004)	-0.023*** (0.006)	- -	- -
Same rig, diff field, diff producer (<i>Experience_{-f-prt}</i>)	- -	- -	-0.005 (0.004)	-0.010 (0.009)
Same rig, same field, diff producer (<i>Experience_{f-prt}</i>)	- -	- -	-0.007 (0.007)	0.001 (0.013)
Same rig, diff field, same producer (<i>Experience_{f-prt}</i>)	- -	- -	-0.011*** (0.004)	-0.006 (0.007)
Same rig, same field, same producer (<i>Experience_{fprt}</i>)	- -	- -	-0.024*** (0.005)	-0.019** (0.009)
Producer X rig fixed effects	N	Y	N	N
Field X producer X rig fixed effects	N	N	N	Y
Number of observations	19059	16325	19059	14289

Parentetical 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, month and year fixed effects, and field, producer, and rig fixed effects

Table 7: Relationship-specific learning results: alternative functional forms
Dependent variable is log(drilling time)

	I	II	III	IV	V
Log of experience with:	Reference Case	Perfect substitution	General CES	Field and producer spline	Total rig experience
Same field, same producer ($Experience_{fpt}$)	-0.030*** (0.006)	-0.031*** (0.006)	-0.031*** (0.006)	spline	-0.030*** (0.006)
Different field, same producer ($Experience_{-fpt}$)	-0.022*** (0.008)	-0.020** (0.008)	-0.020*** (0.006)	spline	-0.022*** (0.008)
Same field, different producer ($Experience_{f-pt}$)	-0.005 (0.005)	-0.005 (0.005)	-0.005 (0.005)	spline	-0.005 (0.005)
Same rig, different producer ($Experience_{-prt}$)	-0.010** (0.004)	-	-	-0.008** (0.004)	-
Same rig ($Experience_{rt}$)	-	-	-	-	-0.018*** (0.007)
Same rig, same producer ($Experience_{prt}$)	-0.025*** (0.004)	-	-	-0.026*** (0.004)	-0.018*** (0.004)
α	-	0.928*** (0.082)	0.712* (0.112)	-	-
ρ	-	-	0.263 (0.432)	-	-
δ	-	-0.033*** (0.008)	-0.040*** (0.007)	-	-
Number of observations	19059	19059	19059	19059	19059

Parentetical values indicate standard errors clustered on field.

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

Specification in columns II and III includes a CES function of rig experience given by $(\delta/\rho) \cdot \log((1-\alpha)E_{-prt}^\rho + \alpha E_{prt}^\rho)$. ρ is set to 1 in column II.

Statistical significance for α is for a test against a null hypothesis of $\alpha=0.5$

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

Table 8: Relationship-specific learning results: alternative specifications
Dependent variable is log(drilling time)

	I	II	III	IV	V
Log of experience with:	Reference case	Stable contract types	Wells 8000 ft and deeper	Drilling time experience	Drilling time experience, Prod-rig FE
Same field, same producer (<i>Experience_{fpt}</i>)	-0.030 ^{***} (0.006)	-0.031 ^{***} (0.010)	-0.024 ^{***} (0.007)	-0.013 ^{***} (0.003)	-0.012 ^{***} (0.003)
Different field, same producer (<i>Experience_{-fpt}</i>)	-0.022 ^{***} (0.008)	0.007 (0.016)	-0.019 ^{**} (0.009)	-0.014 ^{***} (0.003)	-0.002 (0.008)
Same field, different producer (<i>Experience_{f-pt}</i>)	-0.005 (0.005)	0.006 (0.011)	-0.007 (0.007)	-0.003 (0.005)	0.001 (0.004)
Same rig, different producer (<i>Experience_{-prt}</i>)	-0.010 ^{**} (0.004)	9.1E-06 (0.008)	-0.009 ^{**} (0.004)	-0.009 ^{**} (0.004)	-0.003 (0.006)
Same rig, same producer (<i>Experience_{prt}</i>)	-0.025 ^{***} (0.004)	-0.022 ^{***} (0.007)	-0.023 ^{***} (0.005)	-0.027 ^{***} (0.004)	-0.024 ^{***} (0.006)
Producer X rig fixed effects	N	Y	N	N	Y
Number of observations	19059	11877	12581	19059	16325

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, month and year fixed effects, and field, producer, and rig fixed effects

Table 9: Relationship-specific learning results: robustness of experience calculation
Dependent variable is log(drilling time)

	I	II	III	IV
Log of experience with:	Reference Case	One-year experience	Three-year experience	Five-year experience
Same field, same producer (<i>Experience_{fpt}</i>)	-0.030 ^{***} (0.006)	-0.033 ^{***} (0.006)	-0.025 ^{***} (0.006)	-0.026 ^{***} (0.006)
Different field, same producer (<i>Experience_{-fpt}</i>)	-0.022 ^{***} (0.008)	-0.025 ^{***} (0.008)	-0.018 ^{***} (0.006)	-0.014 (0.011)
Same field, different producer (<i>Experience_{f-pt}</i>)	-0.005 (0.005)	-0.006 (0.007)	-0.007 (0.007)	-0.004 (0.008)
Same rig, different producer (<i>Experience_{-prt}</i>)	-0.010 ^{**} (0.004)	-0.010 ^{**} (0.004)	-0.009 ^{**} (0.004)	-0.010 [*] (0.005)
Same rig, same producer (<i>Experience_{prt}</i>)	-0.025 ^{***} (0.004)	-0.025 ^{***} (0.004)	-0.023 ^{***} (0.004)	-0.023 ^{***} (0.004)
Number of observations	19059	19059	17891	15515

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, month and year fixed effects, and field, producer, and rig fixed effects

Table 10: Estimates for the probability a rig is the first to exit its pair
Values shown are marginal effects: $dPr(\text{ExitFirst}) / dX$

	I	II	III	IV	V
	Conditional logit		Linear probability model		
Log of rig's producer-specific experience	-0.061 ^{***} (0.012)	-	-0.059 ^{***} (0.016)	-0.067 ^{**} (0.034)	0.036 (0.093)
Log of rig's total experience	-	-0.008 (0.017)	-	-	-
Pair FE	N/A	N/A	Y	Y	Y
Rig FE	N	N	N	Y	Y
Rig X producer fixed effects	N	N	N	N	Y

Marginal effects calculated at sample mean.

Parenthetical values indicate standard errors clustered on producer.

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