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

Ryan Kellogg

January, 2011

This paper examines learning-by-doing that is specific not just to individual firms but to pairs of firms working together in a contracting relationship. Using data from the oil and gas industry, I find that the productivity of an oil production company and its drilling contractor increases with their joint experience. This learning is relationship-specific: drilling rigs cannot fully appropriate the productivity gains acquired through experience with one production company to their work for another. This result is robust to *ex ante* match specificities. Moreover, producers' and rigs' contracting behavior is consistent with maximization of relationship-specific learning's productivity benefits.

* I am grateful for financial support from the OpenLink Fund at the Coleman Fung Risk Management Center. I thank the editor, Lawrence Katz, and three anonymous referees for suggestions that substantially improved the paper. This paper has also benefited from discussions with Daniel Ackerberg, 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.

I. 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*. 2010). 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 (2007) 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 three 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? 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? Finally, through what mechanisms does relationship-specific learning operate?

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—measured as the number of days required to drill each well—for producers, rigs, and producer-rig pairs.

This paper's 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 use producer-rig fixed effects 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 reduces drilling times by 4.9% on average, yielding cost savings of about $12,400 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 little. I find that this contracting pattern is difficult to explain with switching costs or *ex ante* match specificities alone.

Finally, I examine the mechanism behind the observed relationship-specific learning. While these results are empirically weaker than the primary finding, the learning appears to be driven primarily by the accumulation of personal interactions between the firms' personnel rather than rigs' accumulation of geologic or technical knowledge specific to the wells they drill for their producers. Consistent with this result, I find that firms tend to choose a contractual form that increases cross-firm personal contacts when they anticipate working together for an extended period of time—a situation in which maximizing the rate of relationship-specific learning is particularly important.

While this paper focuses on the oil and gas drilling industry, it seems likely that the prevalence of relationship-specific learning extends beyond the oilpatch. Large construction and manufacturing projects, for example, regularly involve multiple contractors and sub-contractors working under a lead, general contractor. In addition, 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. Becker (1964), Prendergast (1993), and Gibbons and Waldman (2004, 2006) 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 II provides background information on the oil and gas drilling industry, and Section III discusses industry mechanisms for learning-by-doing. Section IV describes the data used in this study. Section V presents a model of learning-by-doing and discusses the empirical strategy. Section VI provides the primary estimation results. Section VII discusses relationship persistence between producers and rigs, and Section VIII examines the mechanisms behind relationship-specific learning. Section IX concludes.

II. Institutional Background

II.A. 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. A lease typically grants a right to operate in only a small part of a field, and most fields are operated by several distinct producers holding different leases.[2]

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

[2] 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.

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.

There is significant variation in geology across fields, particularly with regard to the depth at which they are buried. Some fields are as shallow as 3,000 feet deep 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 knowledge from seismic imaging and previously drilled wells, and are therefore better placed to make these engineering decisions.[3]

II.B. Rigs and Contracting

The actual drilling of wells is conducted by drilling companies, which own drilling rigs and employ drilling crews. The primary features of a rig 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.

[3] Very small producers, which drill infrequently and may not have engineering resources, sometimes outsource the planning and design function to the driller.

Rigs are mobile and can easily change locations within a field; however, moves of more than fifty 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 seven days per week, rotating crews in three eight-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 not known with certainty and contingent on oil production from the first several wells drilled. For 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.[4] To the extent that relationship-specific learning is important, this repeat contracting creates rents that can be bargained over at each renewal.[5] However, unlike classic examples of relationship-specific investments from the transactions cost literature (Williamson 1975, 1985; Klein *et al.* 1978), relationship-specific learning does not generally require costly up-front investment, so the lack of a long-term contract should not create an inefficiency.[6]

[4] 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.

[5] I do not possess proprietary data on specific transactions with which the allocation of rents can be analyzed.

[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. In addition, bargaining over the rents generated by relationship-specific learning should not lead to inefficiencies via premature termination of relationships. In particular, neither firm should ever have an incentive to terminate a relationship in

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 technical specifications regarding the well to be drilled, including for instance the well's total depth, the types of steel casing to be installed in the well, and properties 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 daily operations, typically in consultation with the rig's foreman. The company man has final authority over decisions regarding, for example, which drill bit to use, the weight to place on the drill bit while drilling, and how to respond to drilling problems such as stuck drill pipe.

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. The rig is free to make day-to-day drilling decisions subject to the technical specifications of the RFQ.

Corts and Singh (2004) empirically examine the determinants of contract type in the offshore drilling industry and verify several intuitive theoretical predictions. For example, dayrate contracts are used more frequently in situations in which there is a large amount of geologic uncertainty, since producers are better-equipped with the knowledge and technical staff to address geologic risk. The choice of contractual form also affects drilling performance incentives. Under a footage contract, the rig will have a direct monetary incentive to exert a high level of effort and drill quickly, but this incentive will not exist under a dayrate contract. However, indirect performance incentives are likely to be important. Rig reputations are well-known by producers, and 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

---

order to prevent the other firm from obtaining too much bargaining power: both firms always hold an outside option to find another contracting partner and therefore cannot be made worse off than they would be with a new partner.

in which repeat contracting is common, this observability of effort can generate implicit performance incentives for the contractor (Corts 2007).

III. 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 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. Huckman and Pisano (2006) similarly finds evidence suggesting that the performance of cardiac surgeons, who often work as freelance contractors across multiple hospitals, depends more on their hospital-specific experience than on their general experience.[7]

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 producers accumulate and learn from information with every well they drill. For example, the optimal selection of drilling mud depends critically on the types of rock encountered throughout the well bore, and a producer may drill several wells before discovering the ideal mud composition via an educated trial-and-error process. 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 tends to be less technical in nature than producer-level learning. Instead, rigs' learning comes from improved teamwork and developments in crew members' skills. For example, crews

[7] Neither McCabe (1996) nor Huckman and Pisano (2006) addresses the possibility that the estimated learning effects may be driven by match-specific heterogeneities.

become more efficient at lowering drilling pipe into a hole, ninety 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 geology of its fields, or the producer's company man may improve his knowledge of the capabilities of the rig and its crew. Industry participants have also indicated that repeated personal interactions between the firms' personnel are likely to be important. For instance, the ability to rapidly solve drilling problems, such as the sticking of pipe in the wellbore, is an important determinant of drilling efficiency. These problems may be more easily solved if the company man and rig foreman have, through repeated personal interactions, developed a working relationship that allows them to collaborate effectively.[8]

The intuition behind relationship-specific learning has a parallel in recent theoretical work. Ellison and Holden (2009) 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.

IV. 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

[8] This mechanism suggests that relationship-specific learning occurs between the producer's company man and the rig's 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.

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.[9]

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. Unlike the TRRC data, the unit of observation in the SB data is a rig-week, and I do not observe individual wells. For example, I might see in the SB data that a particular rig worked for 25 consecutive weeks, but these data alone yield no information regarding how many wells it drilled during that time span.

I merge the SB rig location data into the TRRC's drilling records to create a well-level dataset in which each observation reports the well's drilling time, location, producer, and drilling rig. 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: (1) wells drilled in less than one week will not appear in the SB data if there is no drilling on a Friday; (2) names of production companies are not harmonized across the two datasets; (3) the SB data have missing observations; and (4) when a producer drills multiple wells simultaneously in the same county, it is not always possible to match a particular rig to a particular well. Online Appendix 1 discusses the data merge process in detail and notes that match failures are not systematic in a way that could bias the main empirical results.

The data merge yields a dataset with 33,125 observations (wells) 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,

[9] 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 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.

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.[10]

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. Table I indicates that 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 I 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. The summary statistics given at the top of Table II indicate that the sample average drilling time is 23.0 days and the average well depth is 9,036 feet.

V. A Model of Learning-by-Doing in the Drilling Industry

V.A. Production Function

This paper models the objective of producers to be the minimization of the time needed to drill any given well, as motivated by their objective to minimize drilling costs. 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 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 seven days per week, and adding crew members cannot increase the rate of penetration. Most capital drilling inputs, such as the casing and tubing that are installed in the well and the equipment on the rig itself, are fixed functions of the well's depth and the particular rig. 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 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 because the rates charged by rigs rise during periods of high drilling activity, creating spurious correlation between drilling cost and experience.

[10] 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.

My measure of drilling productivity is therefore the rate of drilling in feet per day, which firms try to maximize subject to the constraints imposed by geology, safety, and the physical capabilities of the drilling rig used. Throughout this section, I assume that the producer and rig are aligned in their objective to maximize the rate of drilling given the implicit performance incentives discussed in Section II.B.[11] For a given well—the unit of observation—let y denote the well's inverse rate of drilling in days per foot. y is assumed to be determined by equation (1):

(1) $$y = \phi \cdot g(\Omega) \cdot v$$

ϕ denotes the baseline drilling rate of the rig drilling the well, given its characteristics such as its horsepower. Ω denotes factors such as the efficiency of the rig crew and the decisions the firms make regarding how to drill the well—for example, the choice of drilling mud, the weight to apply onto the drill bit, and the responses to drilling problems such as stuck pipe. The function $g(\Omega)$ translates these factors, which are unobservable, into their effect on the rate of drilling. The firms' objective to maximize the rate of drilling therefore translates to an objective to minimize $g(\Omega)$ by making optimal drilling decisions. That is, firms search for the optimal set of decisions Ω^* that minimizes $g(\Omega)$. Finally, v denotes field and well-specific factors that impact the rate of drilling, such as the characteristics of the rock, the depth of the well within the field (drilling times are superlinear in depth so that drilling rates tend to be lower for deeper wells), and the realization of any drilling problems.[12]

Experience effects—learning-by-doing—do not enter directly into equation (1) but instead affect drilling productivity by improving the quality of the factors Ω, moving them closer to Ω^* and leading to a decrease in $g(\Omega)$. Denote experience by E, and (postponing the definition of E to the next subsection) let $\log(h(E))$ denote the learning process by which experience improves $\log(g(\Omega))$ on average, so that equation (2) denotes a projection of $\log(g(\Omega))$ onto $\log(h(E))$. In (2), $\log(\eta)$ denotes idiosyncratic deviations of $\log(g(\Omega))$ from the expected learning curve given by $\log(h(E))$, potentially reflecting a trial-and-error process of learning. The definition of $\log(h(E))$ as the average learning process implies that $\log(\eta)$ is orthogonal to $\log(h(E))$.

(2) $$\log(g(\Omega)) = \log(h(E)) + \log(\eta)$$

[11] The relationship between contractual form and drilling productivity is explored in Section VIII.C.

[12] Geologic heterogeneity is predominantly cross-field rather than within field. For example, regressing well depth on a set of field fixed effects yields an R^2 of .88.

Let p and r denote the producer and rig that drill the well, f denote the field in which the well is drilled, and t denote the date at which drilling is complete. Allowing for producer-specific "baseline" drilling efficiencies that are independent of experience and for field-specific factors within v, equations (1) and (2) can be combined to yield the fixed effect specification (3), in which γ_f, δ_p, and ϕ_r denote field, producer, and rig fixed effects, respectively:

$$\text{(3)} \qquad \log(y_{fprt}) = \log(h(E)) + \gamma_f + \delta_p + \phi_r + \boldsymbol{\theta X_{fprt}} + \varepsilon_{fprt}$$

The disturbance ε_{fprt} is a composite error term: the sum of $\log(\eta)$ and $\log(v)$. $\boldsymbol{X_{fprt}}$ denotes a vector of observable variables that plausibly impact drilling productivity: (1) a cubic polynomial in the well's depth; (2) dummy variables for whether the well is drilled for oil, gas, or both; (3) dummy variables for month-of-year that capture seasonal weather patterns; and (4) yearly dummy variables that capture any industry-wide technological progress.[13]

V.B. Specification and Calculation of Experience

In the reference case model I let $\log(h(E))$ take the form given by equation (4) below, yielding a log-log specification for learning-by-doing, as is common in the literature:

$$\text{(4)} \qquad \log(h(E)) = \beta_1 \log(E_{ft}) + \beta_2 \log(E_{pt}) + \beta_3 \log(E_{fpt}) + \beta_4 \log(\hat{E}_{rt}) + \beta_5 \log(\hat{E}_{prt})$$

E_{ft} denotes the industry's overall experience in field f on date t, regardless of which firms have been drilling in the field. E_{pt} denotes the overall experience of producer p, and E_{fpt} denotes the experience of producer p in field f. These three variables are analogous to the experience variables used in previous learning-by-doing studies, such as Thornton and Thompson (2001). E_{fpt} measures the experience of the lead firm (p) at a particular type of project (f), while E_{ft} and E_{pt} measure cross-firm and within-firm spillovers, respectively. The primary contribution of this paper is the addition of the $\hat{E}_{rt}$ and $\hat{E}_{prt}$ terms. The former denotes the overall experience of rig r, and the latter denotes the joint experience that rig r and producer p have working together.

The coefficient β_5 measures the strength of relationship-specific learning. For now, the specification is agnostic as to whether relationship-specific learning is driven by personal interactions between the producer and rig or by the rig's increasing familiarity with the technical procedures specific to its producer and/or its producer's fields (task-specific learning, in the

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

language of Gibbons and Waldman [2004, 2006]). This distinction will be explored in Section VIII.B by breaking $\hat{E}_{prt}$ into field-rig, producer-rig, and field-producer-rig specific components.

I measure E_{ft}, E_{pt}, and E_{fpt} as the number of wells recently drilled, while $\hat{E}_{rt}$ and $\hat{E}_{prt}$ are measured as the number of weeks of recent drilling experience (I discuss why I use recent rather than total experience below). The $\hat{E}_{rt}$ and $\hat{E}_{prt}$ variables are measured in units of time rather than wells (E_{rt} and E_{prt}) 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. For the field and producer experience variables, the measures $\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$ that calculate experience as the number of days of recent drilling activity exist in the TRRC data and are potential alternatives to the E_{ft}, E_{pt}, and E_{fpt} measures that count the number of wells drilled. I use E_{ft}, E_{pt}, and E_{fpt} in the reference case empirical specification rather than $\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$ for several reasons. First, the field and producer-specific dimensions of learning tend to be technical and driven by the geologic information gained with each penetration rather than the accumulation of days of experience (a counter-argument would be that, if firms tend to learn more from mistakes than from successes, measuring experience using time spent drilling may be more appropriate). Second, the E_{fpt} measure is used in the petroleum engineering literature, not $\hat{E}_{fpt}$. Third, and finally, the use of E_{ft}, E_{pt}, and E_{fpt} ultimately yields estimates of field-producer learning that are larger in magnitude than when $\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$ are used, consistent with a story in which E_{ft}, E_{pt}, and E_{fpt} are the true measures of experience, which $\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$ measure with error (see columns (1) and (4) of Online Appendix Table A3).

E_{ft}, E_{pt}, and E_{fpt} for a given well are defined as the number of wells for which drilling was completed during the two years prior to date t—the completion date of the well in question.[14] This well itself is also included in E_{ft}, E_{pt}, and E_{fpt}, implying that all wells in the dataset have at least one unit of experience and avoiding a logarithm of zero in a log-log learning specification.[15] I measure experience using drilling within the past two years rather than the total cumulative number of wells drilled because the majority of the fields and firms in the dataset existed prior to

[14] Defining t as the completion date—the date drilling ended—rather than the date drilling began ensures that when the drilling of two wells partially overlaps in time (because multiple rigs are being used simultaneously within a field or by a producer), the well that is completed first counts as experience for the second well. I choose this approach because it allows the second well to benefit from experience gained through drilling the first well, which would have penetrated through any given depth a few days in advance of the second well. If two wells are completed on the same date, neither well counts as experience towards the other.

[15] The "hat" versions of these variables, which are ultimately used as instruments and measure experience as the recently accumulated number of days of drilling, do not include the drilling time of the well completed at t in the experience calculation. Instead, I treat the well completed at t as counting for only 1 day of experience in all cases, thereby ensuring that these experience measures are never zero and enabling the log-log learning specification.

the start of the sample.[16] 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, Beckman, and Epple (1990), Benkard (2000), and Thompson (2007) have demonstrated that experience effects decay with time as learning is "forgotten," supporting the importance of recent experience in determining productivity. Section VI.B discusses evidence of forgetting effects in the drilling industry.

I measure a drilling rig's stand-alone experience $\hat{E}_{rt}$ as the number of weeks the rig was actively drilling within the two years prior to t. Producer-rig joint experience $\hat{E}_{prt}$ is defined similarly (both of these variables also always include a one for the current week). Finally, I calculate all five experience variables using the original TRRC and SB datasets rather than the smaller dataset generated by the match of the TRRC data to the SB data. Were I instead to use this smaller dataset, I would systematically understate each measure of experience. Summary statistics for all five experience measures are presented in Table II.

V.C Estimation and Identification; Producer-Rig Matching

Given the fixed effects and control variables $\boldsymbol{X_{fprt}}$ in equations (3) and (4), the effects of producer and field-level learning (β_1, β_2, and β_3) are identified through variations in each producer's drilling activity within and across each field. There exist numerous sources of such variation, including changes in oil and gas prices (which will have heterogeneous effects across fields), discovery of new fields, and the identification of unexploited reserves in existing fields (through seismic imaging technology, for example). The coefficients on $\hat{E}_{rt}$ and $\hat{E}_{prt}$ are separately 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. Such variation arises because, once a producer and rig begin working together, they do not work with each other indefinitely. Eventually, the producer will run out of wells that it wishes to drill, at least temporarily, and release the rig. For instance, the producer may cease drilling in order to evaluate the oil production from its new wells before continuing with additional drilling. When the producer wishes to re-commence drilling, it may not be able to work with its previous rig if that rig has contracted with another producer, and it will therefore start a relationship with a new rig.

[16] 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 other than two years in Online Appendix 2.

Identification of relationship-specific learning in equations (3) and (4) will fail if rigs have producer-specific characteristics (independently of human capital acquired through learning) and are likely to have longer relationships with producers with which they are well-matched in their ability to efficiently drill wells. Such matching would generate a negative correlation between $\hat{E}_{prt}$ and the disturbance ε_{fprt}, leading to a downward-biased estimate of β_5. As noted in Section II.B, however, a rig's equipment is generally not field or producer-specific, apart from its depth rating, suggesting that match specificities are unlikely to be a serious concern. Nonetheless, I address the potential bias in equation (3) arising from producer-rig specificities by adding fixed effects η_{pr} for producer-rig pairs to the specification, yielding equation (5) below. Equation (5) is the reference case model on which I will focus in estimation.

$$\text{(5)} \qquad \log(y_{fprt}) = \log(h(E)) + \gamma_f + \eta_{pr} + \boldsymbol{\theta X_{fprt}} + \varepsilon_{fprt}$$

With the inclusion of the η_{pr}, identification of relationship-specific learning (β_5) comes only from variation in joint experience within each rig-producer pair.[17] It is important to note that, even with the producer-rig fixed effects included in equation (5), the specification is, strictly speaking, only estimating a learning rate for those producer-rig pairs that are actually observed in the data. If producers are more likely to work with rigs with which they anticipate having steep relationship-specific learning curves, then the hypothetical learning rate for unobserved producer-rig pairs could be lower than the learning rate estimated from equation (5). Short of being able to run a randomized experiment, it is not possible 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, which is the learning rate estimated here.

A second potential impediment to the identification of learning effects is serial correlation in the disturbance term ε_{fprt}, which may present a problem via an argument similar to that made in Benkard's (2000) study of learning in aircraft manufacturing. For the experience variables that are measured in accumulated drilling time—$\hat{E}_{rt}$ and $\hat{E}_{prt}$—positive shocks to ε_{fprt} in the recent past increase the amount of experience. If the disturbances are serially correlated, then $\hat{E}_{rt}$ and $\hat{E}_{prt}$ will be positively correlated with ε_{fprt}, leading to upward bias (towards zero) in the estimated learning rates. That is, serial correlation biases the estimates away from finding evidence of learning. This bias may be particularly important for relationship-specific learning,

[17] β_4 is still identified in the presence of producer-rig fixed effects because rigs sometimes have multiple employment "spells" with a single producer, and E_{rt} will be different in each spell. β_1, β_2, and β_3 are identified when multiple rigs drill simultaneously within a field, producer, and field-producer, respectively.

since it is likely that serial correlation is greatest for wells drilled in sequence by the same producer and rig rather than for wells drilled by different firms. Conversely, for the variables E_{ft}, E_{pt}, and E_{fpt} that calculate experience as the number of wells drilled within the two years prior to date t, the bias from serial correlation works in the opposite direction, exaggerating the estimated learning effect. This bias arises because producers often drill wells one right after the other. In such cases, the number of wells drilled within any fixed time period (such as two years) will be inversely related to the number of days required to drill each well. With serial correlation in the disturbance, positive shocks to the current well's drilling time will be positively correlated with past shocks and therefore negatively correlated with experience, yielding a mechanical downward bias (away from zero) in the estimated learning rates.

Benkard (2000) addressed serial correlation by instrumenting for experience using macroeconomic demand and cost shifters such as the price of oil. This approach is not viable here because, while such shifters can explain changes in aggregate drilling industry experience, they cannot explain firm, relationship, or field-specific experience. Fortunately, it appears that serial correlation in ε_{fprt} is not a substantial issue in this setting: the estimated first-order autocorrelation parameter of the residuals from equation (5), within each producer-rig pair, is only .081.[18] This parameter is much smaller than the values near .5 found in Benkard's (2000) study. Moreover, Benkard found that the interpretation of the results from the instrumental variables estimate was "essentially the same" as that from the non-instrumented estimate, despite the high autocorrelation parameter.[19]

Thus, the impact of serial correlation on the estimated parameters of equation (5) is likely to be minor: the coefficients on $\hat{E}_{rt}$ and $\hat{E}_{prt}$ will be biased slightly upwards (towards zero), and the coefficients on E_{ft}, E_{pt}, and E_{fpt} will be biased slightly downwards. Given that the main parameter of interest is the coefficient on $\hat{E}_{prt}$, the results will therefore be biased against finding evidence of relationship-specific learning. To be conservative in the estimated rates of learning associated with the E_{ft}, E_{pt}, and E_{fpt} terms as well, the reference case estimate of equation (5) instruments for E_{ft}, E_{pt}, and E_{fpt} with measures of experience based on accumulated drilling time: $\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$.[20] Because serial correlation will cause these instruments to be positively

[18] The autocorrelation parameter was estimated following Anderson and Hsiao (1981) to address the small number of observations for some producer-rig pairs. That is, the first-differenced residuals were regressed on lagged first-differenced residuals, using the second lag of the residuals as an instrument.

[19] Serial correlation is relevant only for the most recently produced units (wells, in my case), but the experience variables track production over a long period of time. Thus, serial correlation should cause only a small part of experience to be correlated with the disturbance term, and least-squares estimates should therefore not be substantially biased, consistent with the results from Benkard (2000).

[20] The first stage estimates are strongly statistically significant, as shown in Online Appendix Table A1.

correlated with the disturbance, this approach reverses the direction of bias. As expected, the difference between the reference case and non-instrumented estimates of the coefficients on E_{ft}, E_{pt}, and E_{fpt} is not substantial. In particular, the estimated coefficient on E_{fpt} is -0.023 in the reference case and -0.028 in the non-instrumented estimate (see column (5) of Online Appendix Table A3 for the full set of estimates from the non-instrumented model). The unbiased estimate of this coefficient lies between these two bounds.

VI. Primary Estimation Results for Relationship-Specific Learning

VI.A. Primary Specifications

Table III, column (1), presents the estimated learning rates for the reference case model given by equations (4) and (5). The estimated coefficient on $\log(\hat{E}_{rt})$—the stand-alone experience of the drilling rig—is -0.014 and marginally statistically significant (the p-value is .135). This point estimate implies that a rig with one year's experience will improve its drilling performance, independently of the producer with which it contracts, by 5.4%. The estimated coefficient on $\log(\hat{E}_{prt})$—the joint experience between the rig and the producer to which it is currently contracted—is -0.019 and statistically significant at the 1% level. This estimate implies that maintaining a stable relationship for a full year will yield a relationship-specific productivity improvement of 7.3% that is additional to the 5.4% stand-alone improvement. Thus, on average, rigs with stable contracting relationships are estimated to improve their productivity more than twice as quickly as rigs that frequently change contracting partners.

I use these estimated coefficients to obtain an estimate of the average 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 zero in (4)), the average drilling time in my sample would be increased by 4.9%, equal to 1.13 days at the sample average drilling time of 23.0 days.[21] At the 2005 rig dayrate of approximately \$11,000 per day for a well of average depth, this efficiency gain translates to an average reduction in rig rental cost of approximately \$12,400 per well. Section VII examines the extent to which this cost saving leads firms to maintain long-term relationships.

Other estimated coefficients are also of economic interest. The estimated coefficient on field-producer-specific experience (E_{fpt}) is -0.023 and statistically significant at the 1% level, indicating that producers accumulate field-specific technical expertise as they drill more wells,

[21] The 4.9% figure is equal to the sample average of $100(\exp(-\hat{\beta}_5\hat{E}_{prt})-1)$.

independently of the rigs they hire. There is little evidence of experience spillovers across producers working in the same field or across fields drilled by the same producer: the estimated coefficients on E_{ft} and E_{pt} are negative but small in magnitude and statistically insignificant. This result contrasts with those of Irwin and Klenow (1994) and Thornton and Thompson (2001), which 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 common pool extraction problems: 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.[22] Still, manufacturers of semiconductors should also be unwilling to share knowledge, suggesting that motives for secrecy are unlikely to be the sole driver of the difference in these results. Future research is needed to understand the factors, such as rates of employee turnover, for example, that influence learning spillover rates in different industries.

The estimated coefficients on the control variables $\boldsymbol{X_{fprt}}$ are given in Online Appendix Table A2 and generally agree with intuition. The combination of coefficients on the well depth polynomial indicates that deeper wells have slower drilling rates than shallow wells. The impact of the presence of natural gas is imprecisely estimated. The pattern of month dummies suggests that drilling times tend to be lower in the spring and fall, when temperatures are neither extremely hot nor cold, though the effects are not statistically significant. The year fixed effects suggest the presence of some industry-wide technological improvement in the latter half of the sample, though these coefficients are also not statistically significant.

Column (2) of Table III examines the extent to which match specificities between producers and rigs bias the estimate of relationship-specific learning in the absence of producer-rig fixed effects. This specification estimates the model given by equations (3) and (4), in which there are separate fixed effects for producers and rigs but not fixed effects for producer-rig pairs. The estimated rate of relationship-specific learning is virtually unchanged: the new point estimate on $\log(\hat{E}_{prt})$ is -0.0190 rather than -0.0193. This result is consistent with a limited effect of producer-rig match specificities on relationship durations, consistent with the industry characteristics discussed in Section II.B.[23]

[22] 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.

[23] As a robustness check, I have estimated the reference case specification while including field-producer and field-rig fixed effects in addition to the producer-rig fixed effects. Doing so does not substantially affect the estimated rate of relationship-specific learning: the estimated coefficient on producer-rig experience is -0.017 with a standard error of 0.008. Including field-producer-rig fixed effects yields nearly identical results.

Column (3) of Table III verifies that the finding of substantial relationship-specific learning is not driven by the assumed log-log functional form of equation (4). This specification estimates a flexible functional form in which E_{ft}, E_{pt}, E_{fpt}, and $\hat{E}_{rt}$ enter the specification as flexible splines, while $\hat{E}_{prt}$ continues to enter as a log-log function. That is, I model:

$$\log(h(E)) = s_1(E_{ft}) + s_2(E_{pt}) + s_3(E_{fpt}) + s_4(\hat{E}_{rt}) + \beta_5 \log(\hat{E}_{prt}) \tag{6}$$

where s_1 through s_4 each denotes a 4[th] degree cubic spline. The estimate of β_5 from equations (5) and (6) is -0.018, nearly identical to that of the reference case, confirming that the relationship-specific learning result is not driven by the parameterization in equation (4). I also estimate a version of equation (6) in which all five forms of experience enter through splines. The estimated function $s_5(\hat{E}_{prt})$ is plotted in figure I alongside the log-log learning curve from estimating equation (6). These two curves closely overlie each other, indicating that the log-log relationship is an appropriate choice of functional form.

Online Appendix 2 and Table A3 present additional tests of the robustness of the reference-case results, focusing on alternative measures of experience.

VI.B. Relationship to Previous Learning Studies

To compare the above results with previous work, I estimate the effect of producers' experience on drilling productivity while omitting the influence of their relationships with rigs. This analysis follows previous learning-by-doing studies that investigate lead firm productivity but do not incorporate contractor relationships into the analysis.

The specification given in Table III, column (4) uses only the first three measures of experience from equation (4): E_{ft}, E_{pt}, and E_{fpt}, omitting rig-based experience measures as well as rig and producer-rig fixed effects. In this specification, the estimate of the coefficient on E_{fpt} is -0.038, substantially larger in magnitude that the -0.023 estimated in the reference case, in which rigs were taken into account. Column (5) includes rig fixed effects but continues to omit rig-based experience measures: the estimate of β_3 in this case is -0.032, still greater than the reference case estimate. These results demonstrate that investigations of learning-by-doing that do not consider both the identities of firms' contracting partners and relationship-specific experience may overstate the contribution of lead firms' experience to observed productivity improvements.

Previous work (Argote, Beckman, and Epple 1990; Benkard 2000; Thompson 2007) has also studied institutional forgetting: the decay of experience effects over time. Paralleling these studies, I ask whether drilling experience from the distant past has a smaller effect on current drilling rates than does recent experience. I first examine forgetting effects in a setting in which contracting relationships are not accounted for, following the literature. I focus on the forgetting of field-producer-specific experience E_{fpt} given the lack of evidence of learning spillovers through E_{ft} and E_{pt}. I then examine the forgetting of this experience while taking producers' relationships with rigs' into account, to investigate the extent to which forgetting effects can be explained by relationship-specific learning.

Thus far, I have defined E_{fpt} 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 (7) below, in which $N_{fp\tau}$ denotes the number of wells drilled by producer p in field f on date τ.

$$E_{fpt}(\delta) = \sum_{\tau=t-730}^{t} (e^{\delta(t-\tau)/365} \cdot N_{fp\tau}) \qquad (7)$$

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 equation (7) into equation (8), which does not include measures of experience that involve rigs:

$$\log(y_{fpt}) = \beta_1 \cdot \log(E_{fpt}(\delta)) + \beta_2 \cdot \log(E_{ft}) + \beta_3 \cdot \log(E_{pt}) + \gamma_f + \delta_p + \boldsymbol{\theta X_{fprt}} + \varepsilon_{fpt} \qquad (8)$$

I estimate equation (8) using nonlinear method of moments. As in the non-forgetting estimates, I instrument for the three experience measures using the measures of experience based on accumulated drilling time.[24] Estimation results are shown in column (6) of Table III. I obtain a point estimate of δ equal to -0.928 with a clustered standard error of 0.443, consistent with the presence of forgetting.

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 40% of that made by a well drilled one day ago. This depreciation rate is not as great as that estimated by Argote, Beckman, and Epple (1990) in shipbuilding (for which the corresponding figure is 3.2%), though greater than that estimated by Benkard (2000) in aircraft manufacturing (61%). While this result could reflect literal human forgetting of knowledge or turnover amongst producers' personnel, it may also reflect losses of intellectual capital associated with changes in producers' drilling rigs. I investigate this possibility by augmenting equation (8) with rig fixed effects and variables measuring rig and relationship-specific experience, per equations (3) and

[24] 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.

(4).[25] While the new point estimate of δ from this specification, given in Table III, column (7), is still negative and quite large in magnitude at -2.49, it is estimated imprecisely with a standard error of 1.72. This imprecise estimate suggests that losses of relationship-specific capital between lead firms and contractors may be one of the mechanisms behind the significant estimates of forgetting effects found in other studies. This result also relates to that of Thompson (2007), which found that controlling for labor turnover reduced the estimated rate of forgetting in the shipbuilding industry.

VII. Empirical Analysis of Relationship Persistence

This section empirically examines 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. 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. 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 having less 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 these costs being sunk. This pattern may, however, be consistent with the presence of other *ex ante* match specificities as discussed below.[26]

[25] I do not use producer-rig fixed effects in this specification because field-producer forgetting is not identified in their presence. Field-producer learning itself is only identified when a producer drills multiple wells simultaneously in a field; forgetting effects cannot be identified off of simultaneous drilling activity.

[26] Another alternative explanation is that the pattern reflects matching on the speed of firm-specific (not relationship-specific) learning. That is, a given producer and rig may match because they anticipate fast rates of firm-specific learning with that combination. However, it seems unlikely that such a story is the primary driver of relationship persistence, since relationship-specific learning provides a stronger incentive to maintain a relationship as its benefits are not appropriable across firm pairs. The only incentive to maintain a relationship that is provided by accelerated firm-specific learning comes from the desire to continue learning at a high rate so that future rents will be larger. This incentive will become weak over time as the firm-specific learning curve flattens. In addition, a

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.

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.[27] 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 another producer. I then capture this rig's exit date and the producer-specific experience $\hat{E}_{prt}$ of both rigs at that date. 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 II 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 plotted 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

hazard analysis finds that the hazard rate for relationship termination declines with relationship duration, more in line with the relationship-specific learning interpretation than with the accelerated firm-specific learning interpretation.

[27] 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.

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

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

Estimation results are reported in column (1) of Table IV: 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 with 1 month of experience, the less experienced rig has a 63.7% probability of being the first to exit.[28]

Column (2) of Table IV presents the results of estimating equation (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. First, it may be that some rigs have lower switching costs than others and therefore change jobs frequently. A market in which some rigs are "switchers" and others are "stayers" would generate a LIFO pattern even without relationship-specific learning. Second, *ex ante* match specificities would cause producers to hire rigs with which they match best first and then release those rigs last.

To rule out heterogeneity in rig switching costs, I test for a LIFO pattern while including a set of rig fixed effects in the specification. 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. Column (3) of Table IV indicates that 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. Column (4) demonstrates that, when rig fixed effects are added to the specification, I still find a strong systematic LIFO effect. The estimated marginal effect is -0.067, compared to -0.059 in column (3), and is still statistically significant at the 1% level. Rig heterogeneity is not driving the LIFO result.

[28] Expanding the dataset to use groups with more than two rigs (there are 279 such groups) yields nearly the same point estimate of β as that found for the two-rig groups.

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 1,814 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 (5) of Table IV. The standard error of the estimated marginal effect is nearly five times that of column (3), 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 pre-existing match specificities. I first examine the extent to which there is matching on an observable characteristic of rigs: their depth ratings. For each rig in each pair, I use the absolute difference between the rig's depth rating and the average depth of the wells it drills for the producer as a measure of the quality of the match between the producer and rig. I then examine whether this "depth difference" varies across the first and second rig to enter each pair: a smaller difference for the first rig would be indicative of matching. I find that the distribution across pairs of the first rig's depth difference is very similar to that of the second rig. A Kolmogorov-Smirnov test on these distributions fails to reject equality with a p-value of .565.[29] This result indicates that at least on the basis of observable depth ratings, the two rigs in each pair are equally well-matched to their wells, inconsistent with a matching story.

I also investigate pre-existing performance differences between the rigs in each pair by comparing the drilling times of the first well drilled by each rig. 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 .528.[30] 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, but is instead consistent with a recognition by firms that maintaining long-term relationships helps to maximize the productivity benefits of relationship-specific learning.

[29] The average depth difference for the first rig to enter is 2,870 feet, while that for the second rig is 2,939 feet. A paired t-test fails to reject equality of these means with a t-statistic of 0.56.

[30] This test uses a dataset of 69 pairs for which both rigs can be matched to a TRRC observation (and therefore a drilling time) either upon entry or one week subsequent to entry. I adjust the drilling time of each rig for the effects of the producer's field-specific experience and the rig's overall experience per the estimated reference case learning specification (Table III, column (1)). On average, the adjusted first-well drilling time of the first rig to enter is 2.2% higher than that of the second rig (this difference has a t-statistic of 0.30). Without the experience adjustment, this difference is 2.5%.

VIII. Mechanisms Behind Relationship-Specific Learning

This section examines the mechanisms behind the observed relationship-specific learning. Is this learning driven by repeated personal interactions between the rigs' and producers' personnel or by rigs' increasing familiarity with technical procedures specific to particular producers and their fields? Understanding this learning mechanism is important because it has implications for firms' optimal personnel and contracting practices. For example, if personal interactions are important, firms should have an incentive to maintain cross-firm relationships between their employees by, for example, ensuring that their contracting partners' points of contact within their organizations do not frequently change.

This section presents three sets of results: (1) evidence that relationship-specific learning is driven by a producer's relationship with its specific rig, not the rig's broader drilling company; (2) evidence that rigs' field-specific experience alone is insufficient to drive relationship-specific learning; and (3) evidence that dayrate contracts (in which there is more personal interaction between the producer's and rig's personnel) are associated with a more rapid rate of relationship-specific learning than are footage contracts. No single result is definitive on its own. For example, contract choice is likely to be endogenous, and while the finding of accelerated learning under dayrate contracts is robust to controls for factors that plausibly influence both contract choice and learning, a definitive instrumental variable strategy is not available. The three sets of results are, however, mutually reinforcing in that they are all consistent with a substantial role for personal interactions in driving relationship-specific learning.

VIII.A. Rig Versus Driller Specificity

I first examine whether the importance of producer-rig experience is driven by a given producers' relationships with particular rigs or by their relationships with the drilling companies owning the rigs. I create two new experience variables: $\hat{E}_{dt}$ measures the experience (in weeks of active drilling within the past two years) of drilling company *d*, and $\hat{E}_{pdt}$ measures the joint experience of producer *p* with driller *d*. Column (2) of Table V presents results obtained when these two variables are included in the reference case learning specification (4) and (5). Neither $\hat{E}_{dt}$ nor $\hat{E}_{pdt}$ are estimated to significantly impact drilling times, and the coefficient on $\hat{E}_{prt}$ remains negative and statistically significant, indicating that relationship-specific learning occurs at the producer-rig level rather than the producer-drilling company level. This result implies that this learning is of a nature that cannot easily be transferred across rigs within a drilling company.

In particular, what is being learned by the rigs cannot simply be technical information about the producer's procedures that can easily be written down and shared.

VIII.B. Rigs' Field-Specific Versus Producer-Specific Experience

I next examine whether the estimated producer-rig experience effects are driven by the accumulation of each rig's experience with its current producer, with its current field, or with both. Evidence on this question is already available in columns (2) and (5) of Table III, which compare rates of field and producer-specific learning across specifications that do (column (2)) and do not (column (5)) include rig and relationship-specific experience as covariates. When $\hat{E}_{rt}$ and $\hat{E}_{prt}$ are included, the estimated impacts of field-producer-specific experience E_{fpt} and producer-specific experience E_{pt} are attenuated. However, the estimated impact of field-specific experience E_{ft} in both specifications is essentially identical in magnitude and statistically insignificant. If relationship-specific learning were primarily driven by rigs' field-specific experience, the coefficient on E_{ft} in column (5) should be negative and significant, since rigs would carry field-specific knowledge to different producers working within the same field.[31] The absence of such an effect is consistent with the importance of producer-specificities rather than field-specificities in the mechanism behind the observed relationship-specific learning.

I attempt to further resolve the question of field-specificity vs. producer-specificity by creating two new experience measures. $\hat{E}_{frt}$ measures rig r's experience in field f, regardless of producer, and $\hat{E}_{fprt}$ measures rig r's experience in field f with producer p. 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. Each rig's field-producer-specific experience is calculated similarly. The imputation of these variables suggests that the estimates below should be interpreted with caution as the estimated impacts of $\hat{E}_{frt}$ and $\hat{E}_{fprt}$ may be attenuated by measurement error bias.

Inclusion of $\hat{E}_{frt}$ and $\hat{E}_{fprt}$ in the reference case model yields the estimates presented in column (3) of Table V. The estimated coefficients on $\hat{E}_{prt}$ and $\hat{E}_{frt}$ are small in magnitude and

[31] Within-field, cross-producer rig moves do occur in the data: 5,083 observations are associated with rig-field pairs for which wells are drilled for more than one producer.

statistically insignificant, indicating a lack of evidence that either a rig's producer-specific or field-specific experience on its own is sufficient to significantly improve drilling productivity. However, the sum of the coefficients on $\hat{E}_{prt}$, $\hat{E}_{frt}$, and $\hat{E}_{fprt}$ is equal to -0.025 and statistically significant at the 5% level, providing evidence that accumulating experience within both the same producer and same field is beneficial to productivity.[32] The importance of rigs' producer-specific experience is consistent with the results from columns (2) and (5) of Table III, discussed above, and the importance of their field-specific experience is consistent with producers having different company men and drilling procedures across fields, limiting the cross-field transferability of knowledge. The strength of this result is, however, limited by the fact that it is no longer statistically significant at a conventional test size when fixed effects for the triple interaction of field, producer, and rig identifiers are included in the specification.[33]

VIII.C. Relationship-Specific Learning and Contract Choice

Finally, I examine the difference in learning rates between producers and rigs that work under a dayrate contract and those that work under a footage contract. A differential learning rate between these two contract types can shed light on the mechanism behind relationship-specific learning because under a dayrate contract there is more interaction between producer and rig personnel: the producer always has a company man assigned to the rig and directs the rig's operations in consultation with the rig's foreman. The difference between contracts in direct monetary incentives may also be important, as under a footage contract the rig is the residual claimant and therefore has a particularly strong incentive to improve its productivity.

To evaluate learning rates across different contract types, I include in the regression an interaction of the log of producer-rig-specific experience $\hat{E}_{prt}$ with the fraction α of that experience that was obtained under dayrate contracts. The results of estimating this specification—which for the moment should only be thought of as descriptive rather than

[32] A significant coefficient on the field-producer-rig specific experience term itself is not necessary to draw this conclusion. A test that a rig's field and producer-specific experience are important, not just field-specific experience, requires that the sum of the coefficients on producer-rig and field-producer-rig experience be statistically significant—not the stronger condition that the coefficient on field-producer-rig experience be statistically significant—since adding producer-specific experience within the same field affects both of these terms. The sum of these coefficients is -0.024 and statistically significant at the 5% level. Similarly, the sum of the coefficients on field-rig and field-producer-rig experience is -0.020 and also statistically significant at the 5% level.

[33] These fixed effects address the possibility that the results are driven by match-specificities between fields, producers, and rigs. The fixed effects add 8,069 covariates to the specification and effectively remove from the sample 4,770 observations that are associated with a field-producer-rig triple for which I observe only one well, thereby hindering inference. The sum of the coefficients on the three experience variables is only statistically distinct from zero with a p-value of .232 in this specification.

causal—are presented in column (4) of Table V. The estimated coefficient on $\alpha \cdot \log \hat{E}_{prt}$ is equal to -0.017 and is statistically significant at the 5% level, indicating that relationship-specific learning is stronger under dayrate contracts than under footage contracts.[34] This regression also includes a dummy for whether any given well is drilled under a dayrate contract: the estimated coefficient on this dummy is positive with a point estimate of 0.022 and is marginally statistically significant with a p-value of .171. This result provides some evidence that the initial performance of a producer-rig pair is slightly better with a footage contract than with a dayrate contract; however, this advantage is reversed as experience is accumulated.

A causal interpretation of the coefficient on $\alpha \cdot \log \hat{E}_{prt}$ must take into account the possibility that there exist factors that influence both contract choice and the rate of learning. For example, firms may be more likely to choose a dayrate contract when drilling in fields that present substantial geologic risk, and learning curves in such fields may plausibly be relatively steep. An ideal identification strategy would use an instrumental variable. Unfortunately, such a strategy does not appear to be viable. Candidate instruments, such as measures of firms' size that proxy for their risk appetites, are insufficiently powerful to yield informative results. In the absence of a viable instrumental variable strategy, the next-best approach is to examine the impact of dayrate contracts on learning while controlling for as many observable potentially confounding factors as possible, and it is this approach that I adopt here.

I examine three sets of observables that seem likely to influence the contractual form chosen by a particular producer-rig pair working in a particular field. First, I use well depth, the average drilling time in the field being drilled (across all firms), and the variance of the drilling times in the field being drilled (also across all firms) to measure geologic risk. Second, I use the sizes of the producer and drilling company (as measured by their total number of well-weeks of drilling) to proxy for the firms' willingness to bear risk. Third, I use the experience accumulated by the pair ($\hat{E}_{prt}$) to capture any systematic tendency to switch contractual form over the course of a relationship. The importance of these factors in determining the choice of contract type is examined empirically in Online Appendix 3 and Table A4; the findings from this analysis generally agree with those from Corts and Singh's (2004) study of the offshore drilling industry. In particular, dayrate contracts are more likely when geologic risk is high, when the producer is large, and later in relationships. Corts and Singh attribute this last result to the possibility that

[34] The estimated coefficient of -0.012 (p-value = .118) on the un-interacted producer-rig-specific experience variable indicates that relationship-specific learning is still likely to be occurring under footage contracts, albeit with a relatively low magnitude.

repeated interactions improve trust and thereby reduce the importance of direct performance incentives.

Table V, column (5), examines the rate of learning under dayrate contracts while including interactions of the above determinants of contract choice with $\hat{E}_{prt}$ to control for their potential impact on the rate of learning. $\hat{E}_{prt}$ itself enters this specification as a spline to address the possibility that the rate of learning late in relationships—when dayrate contracts are more likely—may be greater than or less than that given by the log-log functional form. The coefficient on relationship-specific experience accumulated under dayrate contracts, $\alpha \cdot \log \hat{E}_{prt}$, is estimated to be -0.017. This estimate is essentially identical to that from column (4), which omitted these controls, and is statistically significant at the 10% level. The result that learning is accelerated under dayrate contracts is therefore robust to controls for observable, relevant factors, though it may still be biased by unobservables. Column (5) of Table V also presents estimates of the coefficients on the other interaction terms: fields in which the variance of drilling times is particularly high are found to have relatively fast learning rates, but interactions involving other factors such as firm size are not estimated to have a significant effect.

Column (6) of Table V addresses unobservable field-specific factors that could influence both contract choice and learning rates by including interactions of field dummies with relationship-specific experience $\hat{E}_{prt}$. The specification still includes interactions of firm size with $\hat{E}_{prt}$. The point estimate on $\alpha \cdot \log \hat{E}_{prt}$ is -0.016, similar to the results in columns (4) and (5), though it is only statistically distinct from zero with a p-value of .196. This estimate is free of confounds driven by field-specific geology, though it may still be biased if there are unobserved firm-specific factors that affect both learning rates and the choice of contractual form.[35]

The above results suggest that the increase in personal contact under dayrate contracts accelerates the rate of learning and that this effect dominates the direct learning incentives provided by footage contacts. I conclude this section by considering an implication of these results for firms' optimal choice of contractual form. Because the benefits of accelerated learning accrue over time, if firms learn more quickly under dayrate contracts then the likelihood with which they use these contracts should increase with the expected duration of their relationship. I examine whether this systematic pattern holds using the sub-sample of producer-rig pairs that always use the same contractual form. I reduce the data to one observation per field-producer-rig triple, and regress each pair's contract choice on the total duration of their relationship (in logs)

[35] Pushing further and adding interactions of producer and rig dummies with joint experience removes substantial identifying variation, yielding an estimated effect of -0.012 with a standard error of 0.015.

and field, producer, and rig fixed effects.[36] Because the actual relationship duration is equal to the *ex-ante* expected duration plus error, this regression yields an attenuated estimate of the relationship between expected duration and the initial choice of contractual form.

I find that a one-unit change in the log of the relationship duration is associated with a 3.0% increase in the likelihood that a dayrate contract is used. This result is statistically significant at the 1% level.[37] The small magnitude of the estimated effect may reflect both attenuation bias and forces that push in the opposite direction. For example, under a conventional story of contracting costs, footage contracts are more costly to write but provide a stronger productivity incentive, and they should therefore be favored in longer relationships since the initial contracting cost can be spread over a longer period.

This section has presented three sets of results that speak to the mechanism behind relationship-specific learning. While no single set of results is definitive, together all three are consistent with a prominent role for cross-firm personal interactions in driving relationship-specific learning. This finding suggests that the degree of contact between firms' personnel should play a role in operational and contracting decisions, and the evidence that firms' choice of contractual form depends on the expected duration of their relationship is consistent with this view.

IX. 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 producers are more likely to work with rigs with which they have substantial prior experience than those with which they have worked relatively little. Moreover, the observed relationship-

[36] Each producer-rig pair's relationship duration is calculated from the SB data as the difference between the first and last day the pair is observed. This calculation drops the first year of SB data (1990) to avoid pairs that began work prior to the start of the sample. The total number of pairs in the regression is 6,996, and some pairs are observed multiple times due to the field fixed effects. These fixed effects are important because they control for geologic factors that might influence both contract choice and relationship duration. A similar regression that uses only one observation per producer-rig pair and omits field fixed effects has 5,054 observations and yields an estimated effect of 2.5% that is statistically significant at the 1% level.

[37] Column (7) of Table V verifies that the finding of accelerated learning under dayrate contracts from column (5) is robust to the inclusion of an interaction between relationship duration and producer-rig joint experience.

specific learning appears to be driven primarily by the accumulation of personal interactions between the firms' personnel, rather than by just the accumulation of field or firm-specific technical knowledge.

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. Greater technical complexity in an industry's production process, for example, 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. To deepen interpersonal relationships, 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. To the extent that these actions are costly, it may be in the firms' interests to develop contracting arrangements that alleviate *ex post* bargaining problems and promote efficient investment. These personnel and contracting implications of relationship-specific learning are potentially valuable subjects for future research.

DEPARTMENT OF ECONOMICS, UNIVERSITY OF MICHIGAN, AND NBER

References

Ackerberg, Daniel A. and Maristella Botticini, "Endogenous Matching and the Empirical Determinants of Contract Form," *Journal of Political Economy* 110 (2002), 564-591.

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 (2004).

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

Anderson, T.W. and Cheng Hsiao, "Estimation of Dynamic Models with Error Components," *Journal of the American Statistical Association*, 76 (1981), 598-606.

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

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

Arrow, Kenneth J., "The Economic Implications of Learning by Doing," *Review of Economic Studies*, 29 (1962), 155-173.

Becker, Gary S., *Human Capital; a Theoretical and Empirical Analysis, with Special Reference to Education* (New York, NY: Columbia University Press, 1964).

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

——, "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," *Econometrica*, 78 (2010), 453-508.

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

Cabral, Luis M.B. and Michael H. Riordan, "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, 62 (1994), 1115-1140.

Corts, Kenneth S. "The Interaction of Implicit and Explicit Contracts in Construction and Procurement Contracting," University of Toronto working paper (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., 2009).

Gibbons, Robert and Michael Waldman, “Task-Specific Human Capital,” *American Economic Review Papers and Proceedings*, 94 (2004), 203-207.

——, “Enriching a Theory of Wage and Promotion Dynamics Inside Firms,” *Journal of Labor Economics*, 24 (2006), 59-107.

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

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

Jovanovic, Boyan and Yaw Nyarko, “Learning by Doing and the Choice of Technology,” *Econometrica*, 64 (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 (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.

McCabe, Mark J., “Principals, Agents, and the Learning Curve: The Case of Steam-Electric Power Plant Design and Construction,” *Journal of Industrial Economics*, 44 (1996), 357-375.

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

Parente, Stephen L., “Technology Adoption, Learning-by-Doing, and Economic Growth,” *Journal of Economic Theory*, 63 (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 (1993), 523-534.

Spence, Michael A., "The Learning Curve and Competition," *Bell Journal of Economics*, 12 (1981), 49-70.

Stokey, Nancy L., "Learning by Doing and the Introduction of New Goods," *Journal of Political Economy*, 96 (1988), 701-717.

Thompson, Peter, "How Much Did the Liberty Shipbuilders Forget?" *Management Science*, 53 (2007), 908-918.

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 (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 (1985), 368-385.

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

——, *The Economic Institutions of Capitalism* (New York, NY: Free Press, 1985).

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

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

Table I
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

Notes : Data on wells, fields, and producers come from the Texas Railroad Commission (TRRC). Data identifying rigs and drilling companies come from Smith Bits (SB). Data shown include wells drilled from 1991-2005 that successfully merged across the TRRC and SB data. This merge is summarized in Section IV and described in more detail in the Online Appendix.

Table II
Sample Summary Statistics

	Number of observations	Min	Median	Mean	Std. Dev.	Max
Drilling time (days)	19,059	2	18	23.0	19.2	179
Well depth (feet)	19,059	631	9,000	9,036.0	2,817.0	23,000
Gas well (0/1 dummy)	19,059	0	1	0.6	0.5	1
Oil and gas well (0/1 dummy)	19,059	0	0	0.0	0.0	1
Number of wells drilled during the past two years in:						
Same field, same producer	19,059	1	7	22.1	51.3	711
Same field	19,059	1	23	91.5	193.3	1,847
Same producer	19,059	1	66	148.1	184.6	1,108
Number of weeks of drilling within past two years by:						
Same rig	19,059	1	65	60.8	30.5	105
Same rig, same producer	19,059	1	14	27.7	31.0	105

Notes : Data on wells, fields, producers, drilling time, and well characteristics come from the Texas Railroad Commission (TRRC). Data identifying rigs and drilling companies come from Smith Bits (SB). Data shown include wells drilled from 1991-2005 that successfully merged across the TRRC and SB data. This merge is summarized in Section IV and described in more detail in the on-line appendix. See Section V.B for details regarding the construction of the experience variables.

Table III
Learning-By-Doing Estimates. Dependent Variable is log(Drilling Time / Well Depth)

Log of experience with:	(1) Reference case model	(2) No producer-rig fixed effects	(3) Splines	(4) Field variables only	(5) Field variables with Rig FE	(6) Field variables with forgetting	(7) All variables, with forgetting
Same field	-0.009	-0.010	spline	-0.001	-0.010	-0.002	-0.011
(E_{ft})	(0.012)	(0.007)		(0.008)	(0.007)	(0.008)	(0.007)
Same producer	0.003	-0.015	spline	-0.015	-0.024**	-0.016	-0.016
(E_{pt})	(0.020)	(0.011)		(0.011)	(0.011)	(0.011)	(0.011)
Same field, same producer	-0.023***	-0.022***	spline	-0.038***	-0.032***	-0.048***	-0.040***
(E_{fpt})	(0.007)	(0.006)		(0.006)	(0.005)	(0.007)	(0.011)
Same rig	-0.014	-0.019***	spline	-	-	-	-0.020***
($\hat{E}_{rt}$)	(0.009)	(0.006)		-	-	-	(0.006)
Same producer, same rig	-0.019***	-0.019***	-0.018***	-	-	-	-0.016***
($\hat{E}_{prt}$)	(0.007)	(0.004)	(0.006)	-	-	-	(0.004)
Forgetting parameter δ on	-	-	-	-	-	-0.928**	-2.495
E_{fpt}	-	-	-	-	-	(0.443)	(1.718)
Producer fixed effects	Y	Y	Y	Y	Y	Y	Y
Rig fixed effects	Y	Y	Y	N	Y	N	Y
Producer X rig fixed effects	Y	N	Y	N	N	N	N

Notes: Values in parentheses indicate standard errors clustered on producer. *,**,*** indicate significance at the 10%, 5%, and 1% level.

All specifications use 19,059 observations (wells) covering 1991-2005.

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

Measures of field and producer-specific experience E_{ft}, E_{pt}, and E_{fpt} are based on the number of recently drilled wells and are instrumented using the recent number of days spent drilling ($\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$) as discussed in Section V.C. Rig and relationship-specific experience $\hat{E}_{rt}$ and $\hat{E}_{prt}$ are measured as the recent number of weeks spent drilling.

Table IV
Estimates for the Probability that a Rig is the First to Exit its Pair
Values Shown are Marginal Effects: dPr(ExitFirst) / dX

	(1)	(2)	(3)	(4)	(5)
	Conditional logit		Linear probability model		
Log of rig's producer-specific experience ($\hat{E}_{prt}$)	-0.061***	-	-0.059***	-0.067***	0.036
	(0.012)	-	(0.011)	(0.024)	(0.052)
Log of rig's total experience ($\hat{E}_{rt}$)	-	-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

Notes : Marginal effects are calculated at the sample means.

Values in parentheses indicate standard errors clustered on producer.

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

All specifications use 907 "pairs" in which a producer has two rigs drilling wells for it in the same county.

Table V

Empirical Analysis of Learning-By-Doing Mechanisms. Dependent Variable is log(Drilling Time / Well Depth)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log of experience with:	Reference case model	Drilling company experience	Field-producer-rig experience	Contract type	Contract type, controls	Contract type, field FE X Exp	Contract type, duration control
Same field	-0.009	-0.010	-0.008	-0.009	-0.009	0.005	-0.010
(E_{ft})	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)	(0.011)	(0.012)
Same producer	0.003	0.005	-0.002	0.001	0.009	0.005	0.010
(E_{pt})	(0.020)	(0.022)	(0.020)	(0.020)	(0.021)	(0.023)	(0.021)
Same field, same producer	-0.023***	-0.022***	-0.015*	-0.023***	-0.024***	-0.027***	-0.024***
(E_{fpt})	(0.007)	(0.007)	(0.008)	(0.007)	(0.007)	(0.007)	(0.007)
Same rig	-0.014	-0.019*	-0.013	-0.012	-0.012	-0.022**	-0.011
($\hat{E}_{rt}$)	(0.009)	(0.010)	(0.012)	(0.009)	(0.009)	(0.011)	(0.009)
Same producer, same rig	-0.019***	-0.017*	-0.004	-0.012	spline	-	spline
($\hat{E}_{prt}$)	(0.007)	(0.009)	(0.012)	(0.007)		-	
Same field, same rig	-	-	-0.001	-	-	-	-
($\hat{E}_{frt}$)	-	-	(0.015)	-	-	-	-
Same field, same producer,	-	-	-0.019	-	-	-	-
same rig ($\hat{E}_{fprt}$)	-	-	(0.016)	-	-	-	-
Same drilling company	-	0.018	-	-	-	-	-
($\hat{E}_{dt}$)	-	(0.015)	-	-	-	-	-
Same producer, same drilling	-	-0.005	-	-	-	-	-
company ($\hat{E}_{pdt}$)	-	(0.010)	-	-	-	-	-
Dayrate contract dummy	-	-	-	0.022	0.022	0.028	0.022
	-	-	-	(0.016)	(0.016)	(0.017)	(0.016)
Interaction of log($\hat{E}_{prt}$) with:							
Fraction of work on dayrate	-	-	-	-0.017**	-0.017*	-0.016	-0.016*
(α)	-	-	-	(0.008)	(0.010)	(0.012)	(0.010)
Field's average log(drilltime)	-	-	-	-	0.007	-	0.007
	-	-	-	-	(0.011)	-	(0.011)
Field's std. dev. of	-	-	-	-	-0.090***	-	-0.089***
log(drilltime)	-	-	-	-	(0.033)	-	(0.034)
Log(well depth)	-	-	-	-	-0.004	-	-0.003
	-	-	-	-	(0.008)	-	(0.008)
Log(producer size)	-	-	-	-	0.004	0.006	0.005
	-	-	-	-	(0.004)	(0.005)	(0.004)
Log(driller size)	-	-	-	-	0.001	-8.4E-05	0.002
	-	-	-	-	(0.003)	(0.004)	(0.004)
Log(total duration of	-	-	-	-	-	-	-0.009
relationship)	-	-	-	-	-	-	(0.006)
Interactions of field fixed effects with $\hat{E}_{prt}$	N	N	N	N	N	Y	N

Notes: Values in parentheses indicate standard errors clustered on producer. *,**,*** indicate significance at the 10%, 5%, and 1% level.

All specifications use 19,059 observations (wells) covering 1991-2005.

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

Measures of field and producer-specific experience E_{ft}, E_{pt}, and E_{fpt} are based on the number of recently drilled wells and are instrumented using the recent number of days spent drilling ($\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$) as discussed in Section V.C. Rig and relationship-specific experience $\hat{E}_{rt}$ and $\hat{E}_{prt}$ are measured as the recent number of weeks spent drilling.

Figure I
Estimated Spline for Relationship-Specific Learning

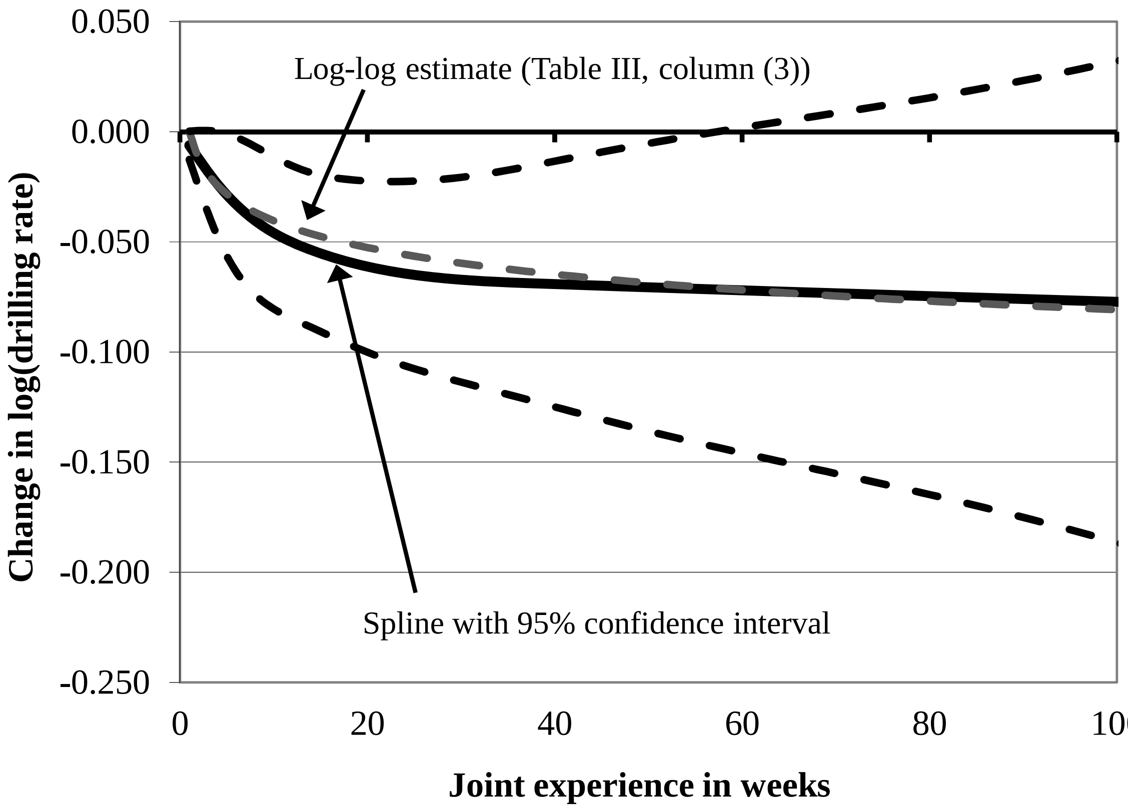

Both the spline and log-log estimates of the impact of joint producer-rig experience ($\hat{E}_{prt}$) are from a model in which all other forms of experience are modeled as splines.

Figure II
Likelihood that the Least Experienced Rig is the First to Change Producers
vs. the Within-Pair Difference in Rigs' Producer-Specific Experience

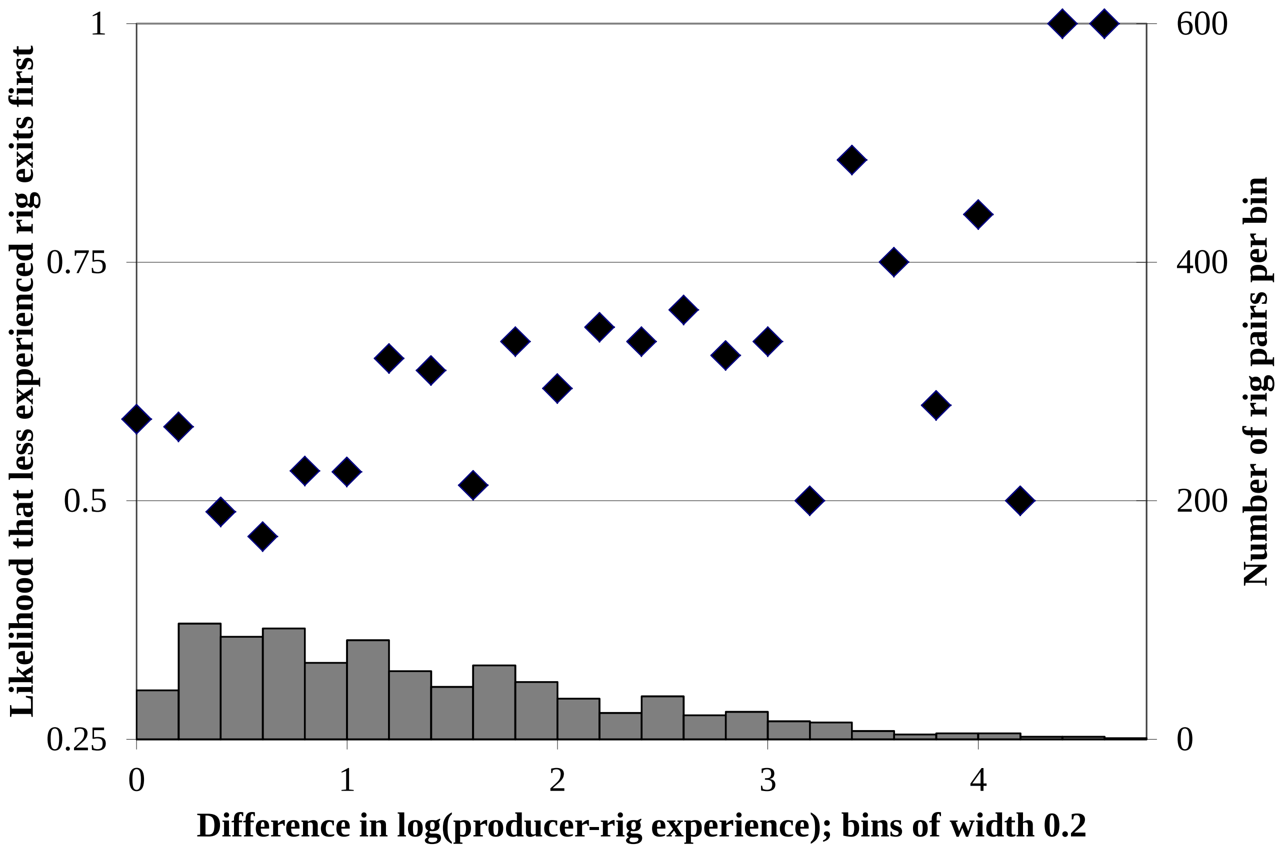

Online Appendices for "Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch"

by Ryan Kellogg

Appendix 1

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 have no corresponding record in the SB data and are therefore impossible to match. Although these wells constitute only 6 percent of the overall population, it is possible that this selection on the dependent variable may bias the results. I address this concern by estimating (4) and (5) with data only for wells that are at least 8000 feet deep (12,128 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 (2) of table A3, are similar to those obtained from the full sample, shown in column (1).

Second, 14.8 percent 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 percent 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 percent 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. 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-by-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.

Finally, some non-unique matches occur when a producer employs multiple drilling rigs simultaneously in the same county. Because the SB data do not contain field or well information, I am unable to distinguish which rig is drilling which well in such cases. While I am able to use information on well depth and well type to match some of these wells to their rigs, there are other cases in which there is no way to confidently match the data. I drop all wells that cannot be matched uniquely, reducing the dataset by a further 20.4 percent of the original TRRC well count.

Appendix 2

This appendix verifies the robustness of the reference case relationship-specific learning result (column (1) of tables III and A3) to a series of alternative specifications and variable definitions, the results of which are reported in table A3.[1] Column (3) considers a specification in which $\hat{E}_{rt}$ is replaced by $\hat{E}_{-prt}$, which measures the experience of rig r with producers other than p. In this regression, the importance of relationship-specific learning is given by the difference between the coefficients on $\hat{E}_{-prt}$ and $\hat{E}_{prt}$. I estimate this difference to be equal to -0.021, very similar to the -0.019 estimate on $\hat{E}_{prt}$ in the reference case and statistically significant with a p-value of .019.

The results reported in the main text measure field and producer experience using E_{ft}, E_{pt}, and E_{fpt}: the number of wells drilled within two years of date t. Table A3, column (4), presents the results of estimating an alternative specification in which these experience variables are measured directly by $\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$: the number of days of drilling within two years of date t. The estimated relationship-specific learning effect is not substantially affected by this change. The importance of the producer's field-specific experience, however, is diminished relative to the reference case model. The source of this change may be that, if E_{fpt} is a better reflection of the process by which learning occurs, then $\hat{E}_{fpt}$ measures experience with error and the point estimate in column (4) reflects attenuation bias.

Column (5) estimates the model without instrumenting for E_{ft}, E_{pt}, and E_{fpt} with $\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$. As expected, the difference between the reference case and non-instrumented estimates of the coefficients on E_{ft}, E_{pt}, and E_{fpt} are not substantial. These results reinforce the intuition that the bias generated by serial correlation in the disturbance ε_{fprt} is minor. Moreover, the estimated

[1] For additional specifications and robustness tests, see Kellogg, Ryan, "Learning by Drilling: Inter-Firm Learning and Relationship Persistence in the Texas Oilpatch," NBER Working Paper #15060.

relationship-specific learning effect is not substantially affected by using a least squares rather than instrumental variables estimator.

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 (6), (7), and (8) calculate experience over one, three, and five years, respectively, rather than the reference case of two years. At the one-year horizon, the point estimate on $\hat{E}_{prt}$ is reduced to -0.011, though it is still statistically significant with a p-value of .057. At three and five year horizons, the estimated relationship-specific learning rate is actually stronger than that estimated in the reference case.[2]

Appendix 3

This appendix examines the determinants of whether producers and rigs choose to work together under dayrate or footage contracts. I focus on three sets of factors: those related to the geologic risk and difficulty of drilling in a particular field, those related to the risk appetite of each firm, and those related to the number of previous interactions between the producer and the rig. Factors relating to geologic risk are measured using three variables: the average (across all producers) of log(drilling time) across all wells drilled in the field, the standard deviation (across all producers) of log(drilling time) across all wells drilled in the field, and the depth of the well being drilled. The firms' risk appetites are measured using proxies for firm size: the log of the total number of well-weeks of drilling observed in the sample by each producer and drilling company. Finally, I measure previous interactions using the log of each producer-rig pair's joint experience: $\hat{E}_{prt}$.

For each drilled well in the sample, I regress the binary choice of contract on all six variables and report the estimation results in column (1) of table A4. As with the learning regressions in the main text, standard errors are clustered on producer to address heteroscedasticity and serial correlation. The three measures of geologic difficulty and risk are all estimated to be associated with a greater frequency of dayrate contracting, as expected. Relatively large producers are estimated to be more likely to use dayrate contracts, and the likelihood of dayrate contracting is also found to increase with the number of previous interactions between the producer and the rig. All of these results are aligned with the findings of Corts and Singh (2004) for offshore drilling in the Gulf of Mexico. The one finding that diverges from that paper is that relatively large drilling companies tend to also favor dayrate contracts. A

[2] The number of observations for these estimates is reduced because calculating experience over more than two years requires dropping observations that occur early in the sample.

theory of risk allocation would predict the opposite effect. One story that could explain this result is that relatively large firms have well-known reputations for being efficient drillers; thus, producers do not need to use footage contracts when working with them. Instead, they use dayrate contracts that are less costly to write and induce faster learning.

The estimates in table A4, column (1), must be interpreted with caution because they are subject to the criticism of Ackerberg and Botticini (2002). If there are unobserved characteristics of firms (such as the portion of their risk tolerance that is not captured by the firm size variables) that are correlated both with their choice of contract type and field in which to work, then the estimated coefficients on the variables measuring geologic risk may be biased. Similarly, if fields have unobserved characteristics that are not captured by the observed variables, then the estimated coefficients on the firm size variables that proxy for risk appetites may be biased.

To address these concerns, I use fixed effects estimators that control for these unobserved characteristics, under the assumption that the characteristics are constant over the sample. Table A4, column (2), estimates the influence of firm size on contract choice in the presence of field fixed effects, and still finds that both large producers and large drillers are more likely to choose dayrate contracts. The finding that dayrate contracts are more likely late in a relationship is also robust to these fixed effects. Column (3) replaces the field fixed effects with producer-driller fixed effects. Here, I still estimate that relatively deep wells are more likely to be drilled under a dayrate contract—the estimated coefficient actually increases substantially in magnitude, from 0.022 in column (1) to 0.186 in column (3). Because field depth is strongly correlated with drilling times, this increase in the depth coefficient leads the estimated coefficient on fields' average drilling times to be negative and marginally statistically significant with a p-value of .128. The coefficient on the standard deviation of drilling times is still positive, though attenuated in magnitude and statistically distinct from zero with a p-value of only .158. Repeated contracting is still estimated to significantly increase the likelihood that a dayrate contract is used. Overall, these fixed effects estimates indicate that the baseline results from column (1) were not wholly driven by matching on unobserved characteristics, though accounting for this matching does affect which particular field-level variables are estimated to be the most important drivers of contractual form.

Table A1

First stage regressions for field and producer experience variables E_{ft}, E_{pt}, and E_{fpt}

Log of experience based on drilling time:	(1) Same field (E_{ft})	(2) Same producer (E_{pt})	(3) Same field, same producer (E_{fpt})
	Dependent variable: log of experience based on well counts		
Same field ($\hat{E}_{ft}$)	0.511***	0.002	0.013
	(0.012)	(0.006)	(0.012)
Same producer ($\hat{E}_{pt}$)	0.049***	0.685***	0.087***
	(0.015)	(0.031)	(0.023)
Same field, same producer ($\hat{E}_{fpt}$)	0.012***	0.006**	0.439***
	(0.004)	(0.003)	(0.010)
Same rig ($\hat{E}_{rt}$)	0.012	-0.009	-0.007
	(0.009)	(0.009)	(0.013)
Same producer, same rig ($\hat{E}_{prt}$)	0.020***	0.039***	0.065***
	(0.005)	(0.006)	(0.009)
Number of observations	19059	19059	19059

Notes : Values in parentheses indicate standard errors clustered on producer. *,**,*** indicate significance at the 10%, 5%, and 1% level.

All specifications use 19,059 observations (wells) covering 1991-2005.

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

Table A2

Learning-by-doing estimates for the reference case model (table IV, column (1)), including control variables. Dependent variable is log(drilling time / well depth)

Variable	Point estimate	Standard error		Variable	Point estimate	Standard error
Log of experience with:				September dummy	0.017	(0.014)
Same field (E_{ft})	-0.009	(0.012)		October dummy	-0.016	(0.016)
Same producer (E_{pt})	0.003	(0.020)		November dummy	-0.010	(0.017)
Same field, same producer (E_{fpt})	-0.023	(0.007)	***	December dummy	-0.013	(0.019)
Same rig ($\hat{E}_{rt}$)	-0.014	(0.009)		1992 dummy	-0.037	(0.037)
Same producer, same rig ($\hat{E}_{prt}$)	-0.019	(0.007)	***	1993 dummy	0.001	(0.037)
Well depth / 1000	-0.346	(0.144)	**	1994 dummy	-0.020	(0.045)
(Well depth)^2 / 1,000,000	0.034	(0.013)	***	1995 dummy	-0.012	(0.047)
(Well depth)^3 / 1,000,000,000	-8.9E-04	-3.5E-04	**	1996 dummy	0.012	(0.052)
Gas well	3.3E-05	(0.025)		1997 dummy	-0.004	(0.058)
Oil and gas well	0.201	(0.118)	*	1998 dummy	0.024	(0.065)
February dummy	-6.1E-04	(0.016)		1999 dummy	0.019	(0.073)
March dummy	-0.017	(0.016)		2000 dummy	-0.033	(0.070)
April dummy	-0.017	(0.018)		2001 dummy	-0.010	(0.075)
May dummy	-0.027	(0.017)		2002 dummy	-0.024	(0.078)
June dummy	-0.032	(0.018)	*	2003 dummy	-0.060	(0.081)
July dummy	-0.018	(0.017)		2004 dummy	-0.099	(0.082)
August dummy	-0.004	(0.015)		2005 dummy	-0.099	(0.088)

Notes: Values in parentheses indicate standard errors clustered on producer. *,**,*** indicate significance at the 10%, 5%, and 1% level.

Specification uses 19,059 observations (wells) covering 1991-2005.

Regression includes field fixed effects and producer X rig fixed effects.

Measures of field and producer-specific experience E_{ft}, E_{pt}, and E_{fpt} are based on the number of recently drilled wells and are instrumented using the recent number of days spent drilling ($\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$) as discussed in Section V.C. Rig and relationship-specific experience $\hat{E}_{rt}$ and $\hat{E}_{prt}$ are measured as the recent number of weeks spent drilling.

Table A3

Empirical analysis of learning-by-doing: robustness. Dependent variable is log(drilling time / well depth)

Log of experience with:	(1) Reference case model	(2) Wells deeper than 8000 feet	(3) Same rig, other producer experience	(4) Drilltime-based experience	(5) No IV	(6) One-year experience	(7) Three-year experience	(8) Five-year experience
Same field	-0.009	-0.026*	-0.010	-0.005	-0.005	-0.016	-0.013	-0.025
(E_{ft})	(0.012)	(0.015)	(0.012)	(0.006)	(0.011)	(0.011)	(0.013)	(0.017)
Same producer	0.003	0.019	0.003	-0.001	-0.005	-0.011	0.019	0.031
(E_{pt})	(0.020)	(0.021)	(0.020)	(0.014)	0.017	(0.016)	(0.023)	(0.024)
Same field, same producer	-0.023***	-0.018**	-0.022***	-0.010***	-0.028***	-0.025***	-0.018**	-0.019***
(E_{fpt})	(0.007)	(0.008)	(0.007)	(0.003)	(0.007)	(0.007)	(0.008)	(0.007)
Same rig	-0.014	-0.019*	-	-0.014	-0.015	-0.014	-0.013	-0.010
($\hat{E}_{rt}$)	(0.009)	(0.012)	-	(0.009)	(0.009)	(0.009)	(0.010)	(0.011)
Same producer, same rig	-0.019***	-0.017**	-0.024***	-0.021***	-0.017***	-0.011*	-0.023***	-0.027***
($\hat{E}_{prt}$)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
Other producer, same rig	-	-	-0.004	-	-	-	-	-
($\hat{E}_{-prt}$)	-	-	(0.005)	-	-	-	-	-
Number of observations	19059	12128	19059	19059	19059	19059	17891	15515

Notes : Values in parentheses indicate standard errors clustered on producer. *,**,*** indicate significance at the 10%, 5%, and 1% level.

All specifications use 19,059 observations (wells) covering 1991-2005.

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

Measures of field and producer-specific experience E_{ft}, E_{pt}, and E_{fpt} are based on the number of recently drilled wells and are instrumented using the recent number of days spent drilling ($\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$) as discussed in Section V.C. Exceptions are column (4), in which $\hat{E}_{ft}$, $\hat{E}_{pt}$, and $\hat{E}_{fpt}$ are used directly as covariates, and column (5), in which E_{ft}, E_{pt}, and E_{fpt} are not instrumented. Rig and relationship-specific experience $\hat{E}_{rt}$ and $\hat{E}_{prt}$ are measured as the recent number of weeks spent drilling.

Table A4
Empirical analysis of contract choice
Dependent variable is binary: 0 for footage contract, 1 for dayrate

	(1)	(2)	(3)
	No fixed effects	Field fixed effects	Producer and driller fixed effects
Mean log drilling time within field	0.075**	-	-0.037
	(0.037)	-	(0.024)
Std. deviation of log drilling time within field	0.168***	-	0.045
	(0.056)	-	(0.032)
Log(well depth)	0.022***	-	0.186**
	(0.059)	-	(0.040)
Log(producer size)	0.048***	0.049**	-
	(0.013)	(0.009)	-
Log(driller size)	0.087***	0.076***	-
	(0.009)	(0.008)	-
Log of experience within same producer and rig ($\hat{E}_{prt}$)	0.032***	0.031***	0.024***
	(0.008)	(0.005)	(0.004)
Number of observations	19059	19059	19059

Notes : Values in parentheses indicate standard errors clustered on producer.

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

All specifications use 19,059 observations (wells) covering 1991-2005.